arXiv:1608.00948v1 [stat.ME] 2 Aug 2016

The bootstrap, covariance matrices and PCA in moderate and high-dimensions

Noureddine El Karoui*
and
Elizabeth Purdom
Department of Statistics, University of California, Berkeley

June 18, 2016

Abstract

We consider the properties of the bootstrap as a tool for inference concerning the eigenvalues of a sample covariance matrix computed from an $n \times p$ data matrix X. We focus on the modern framework where p/n is not close to 0 but remains bounded as n and p tend to infinity.

Through a mix of numerical and theoretical considerations, we show that the bootstrap is not in general a reliable inferential tool in the setting we consider. However, in the case where the population covariance matrix is well-approximated by a finite rank matrix, the bootstrap performs as it does in finite dimension.

1 Introduction

The bootstrap [18] is a central tool of applied statistics, enabling inference by assessing the variability of the statistics of interest directly from the data and without explicit appeal to asymptotic theory. The appeal of the bootstrap is especially great when asymptotic theoretical derivations are difficult and/or can be done only under quite restrictive assumptions. For instance, consider the case of Principal Components Analysis (PCA). The classic text of [2, 3] (Chapter 13) gives limit theory for the eigenvalues and eigenvectors of the sample covariance matrix when the data is drawn from a normal population. These limit results are non-trivial to derive, even in the Gaussian case, and depend, for instance, on assumptions regarding the multiplicity of the eigenvalues of the population covariance matrix. Furthermore, it is clear, using approximation arguments from [33], that these limit results are not valid for a broad class of distributions. For instance, they do not apply to populations distributions with kurtosis not equal to 3. The modern theory of PCA which aims for better finite-sample approximations by relaxing the assumption that $p/n \to 0$ is much more difficult technically and relies on very strong assumptions about the geometry of the dataset (see [32, 31, 21], follow-up papers, and Section 1.2 below for a short summary).

Remarkably, from a theoretical standpoint, it has been shown that in many situations the bootstrap estimates the distribution of the statistics of interest accurately, at least with sufficient sample

*The authors gratefully acknowledge the support of grant NSF DMS-1510172. N. El Karoui thanks Profs. H-T Yau and Michael Brenner for discussions while visiting Harvard's Center of Mathematical Sciences and Applications; **Keywords**: bootstrap, principal component analysis, random matrices. **AMS 2010 MSC**: Primary 62G09; Secondary: 62H25

sizes (see [10, 27] for classic references). For the specific example of estimating the eigenvalues of the sample covariance matrix and PCA, numerous papers have been written about the properties of the bootstrap [7, 1, 17, 16, 28]. The main results of these papers is that the bootstrap works in an asymptotic regime that assumes that the sample size grows to infinity while the dimension of the data is fixed, with the additional provision that the population covariance has eigenvalues of multiplicity one. When the assumption of multiplicity equal to one does not hold, subsampling techniques [43] can be used to correctly estimate the distributions of interest by resampling. We note however that these subsampling techniques also require the statistician to have subsamples of size that is infinitely large compared to the dimension of the data.

Given the limitations of existing asymptotic theory and these theoretical results on bootstrapping of eigenvalues, it is not surprising that the bootstrap is a natural tool to use in connection with PCA and inferential questions therein. The bootstrap is mostly used in this context to assess variability of eigenvalues, for instance to come up with principled cutoff selections in PCA and related methods such as factor analysis. For recent examples of an applied nature, we refer the reader to [49, 23, 48, 4]. Another application of the bootstrap is of course in bagging [12]; a well known instance of bagging related to high-dimensional covariance estimation is in resampled portfolio selection [37].

Our framework: p/n not close to zero The theoretical assumptions that support the use of the bootstrap make the fundamental assumption that the dimension p is much smaller than n (i.e. $p/n \to 0$). The modern asymptotic theory of PCA, referenced above, has shown that relaxing that assumption – for example by assuming that $p/n < 1$ but does not tend to 0 – leads to dramatically different theoretical behavior of the eigenvalues and eigenvectors. This leads us to question what effect this assumption has on the performance of the bootstrap in the situation where p/n is not close to 0.

In addition its theoretical interest, the asymptotic analysis in this framework tends to yield very accurate finite-sample approximations[32] and hence gives accurate information concerning the practical performance of the methods we consider. Furthermore, in statistical practice p/n is rarely very close to 0, and hence classical approximations, which rely heavily on that assumption, may lead to theoretical results and interpretations that differ quite drastically from what is observed by practitioners.

However, when p/n is not close to 0, developing theoretical results is still quite technically difficult. It requires a large variety of tools, whether one is concerned with the properties of the bulk of the eigenvalues ([36, 50, 46]), or the largest ones ([32, 20, 35]) – which are particularly important in Principal Component Analysis (PCA). These considerations motivate our exploration of the bootstrap as an alternative, data-driven way, to perform inferential tasks for spectral properties of large covariance matrices.

Contributions of the paper The paper is divided into two main sections. In Section 2, we study the performance of the bootstrap by simulations in the context of PCA. We assess whether the bootstrap recovers the sampling distribution of various statistics of interest, for instance the largest eigenvalue of a covariance matrix. We also consider the simpler problem of whether the bootstrap estimates of bias and variance can be used to accurately measure the bias and the variance of various statistics. Most of our results are negative. Only when the largest eigenvalues become quite large compared to the rest does the bootstrap provide accurate inference. Furthermore, the behavior of the bootstrap is very unpredictable: for instance, in two setups that are nearly similar from a population

standpoint, the bootstrap estimate of bias is itself biased, but in one case it underestimates the true bias and in the other it overestimates it.

In Section 3 we provide theoretical results that help explain this behavior. Those results concern two different aspects of the bootstrap. The first results are about the behavior of the bootstrapped empirical distribution of *all* the eigenvalues of the sample covariance matrix $\widehat{\Sigma}$. We show that in the framework we consider (p/n not very close to zero) the bootstrapped empirical distribution is biased and asymptotically non-random. We then consider the bootstrap behavior of only the largest eigenvalues of $\widehat{\Sigma}$. We show that when the population covariance Σ has some very large eigenvalues, far separated from the other eigenvalues, the bootstrap distribution of those large eigenvalues correctly approximates the sampling distribution of the large eigenvalues of $\widehat{\Sigma}$.

The results of this paper confirm that the bootstrap works when the problem is very low-dimensional or can be approximated by a very low-dimensional problem, but is untrustworthy when the problem is genuinely high-dimensional. As such, the current paper complements the findings of the paper [22] that was concerned with the bootstrap for linear regression models.

We now give basic notation and background regarding the bootstrap and estimation of covariance matrices in high dimensions.

1.1 Notations and default conventions

If X is an $n \times p$ data matrix, we call $\widehat{\Sigma}$ its associated covariance matrix, i.e

$$\widehat{\Sigma} = \frac{1}{n-1}(X - \bar{X})'(X - \bar{X}) \ .$$

We also use the notation $\tilde{X} \triangleq (X - \bar{X})$.

We call empirical spectral distribution of a $p \times p$ symmetric matrix M the probability measure such that $dF_p(x) = \frac{1}{p}\sum_{i=1}^{p} \delta_{\lambda_i(M)}$, where $\lambda_1(M) \geq \lambda_2(M) \geq \ldots \geq \lambda_p(M)$ are the ordered eigenvalues of M. We also use the notation $\lambda_{\max}(M)$ for the largest eigenvalue of the matrix M.

For $z \in \mathbb{C}^+$, i.e $z = u + iv$, where $v > 0$, we call $m_p(z)$ the Stieltjes transform of the distribution F_p, i.e

$$m_p(z) = \int \frac{1}{x - z} dF_p(x) = \frac{1}{p}\text{trace}\left((\widehat{\Sigma} - z\text{Id}_p)^{-1}\right) \ .$$

We use the notation $\Longrightarrow$ to denote weak convergence of probability distributions.

$|||M|||_2$ is the operator norm of M, i.e its largest singular value. $\|w\|_\infty = \max_{1 \leq i \leq p} |w_i|$ is the ℓ_∞-norm of the vector w. We say that the sequence $u_n = \text{polyLog}(n)$ if u_n grows at most like a polynomial in $\log(n)$.

In doing asymptotic analysis, we work under the assumptions that $p/n \to r$, $r \in (0, \infty)$.

We call the *Gaussian phase transition* the value $1 + \sqrt{p/n}$ (see the comments after Theorem 1.1 for details).

1.2 Review of Theoretical results about High-dimensional Covariance Matrices

We review briefly two key results concerning the spectral properties of high-dimensional covariance matrices. More details and more general background results are in the Supplementary Material.

The first result (from [32]) concerns the distribution of the largest eigenvalue of $\widehat{\Sigma}$ in the case that might be considered the "null" case for PCA – the predictors are all independent with covariance matrix equal to Id_p.

Theorem 1.1 (Johnstone, '01)**.** *Suppose the design matrix X has $X_{i,j} \overset{iid}{\backsim} \mathcal{N}(0,1)$. Call $\widehat{\Sigma}$ the sample covariance matrix of X. Let $\lambda_1(\widehat{\Sigma})$ be the largest eigenvalue of $\widehat{\Sigma}$. Then, if $p/n \to r$, $r \in (0, \infty)$,*

$$n^{2/3}\frac{\lambda_1(\widehat{\Sigma}) - \mu_{n,p}}{\sigma_{n,p}} \Longrightarrow TW_1 \ ,$$
$$\mu_{n,p} \to (1+\sqrt{r})^2$$

TW_1 refers to the Tracy-Widom distribution appearing in the study of the Gaussian Orthogonal ensemble (GOE); details about its density can be found in [32] for instance. Further details about $\mu_{n,p}$ and $\sigma_{n,p}$ are in the appendix; they both converge to finite, non-zero limit. This result implies, among other things, that the standard estimate $\lambda_1(\widehat{\Sigma})$ is a biased estimator of the true $\lambda_1(\Sigma)$ when p/n is not close to zero, overestimating the true size of $\lambda_1(\Sigma)$.

From the point of view of PCA, $\lambda_1(\Sigma) > 1$ corresponds to the scenario of the "alternative" hypothesis, and an important question is how well can we differentiate when the data came from the alternative distribution rather than the null. [6] shows that the distribution of $\lambda_1(\widehat{\Sigma})$ given in Theorem 1.1 also describes the distribution of $\lambda_1(\widehat{\Sigma})$ when Σ is a finite rank perturbations of the Id_p, provided none of the eigenvalues of Σ are too separated from each other. In practical terms, this signifies that $\lambda_1(\widehat{\Sigma})$ has - asymptotically - the exact same distribution under the null as under the alternative and therefore no ability to differentiate the null and the alternative, provided the alternative is not far away from the null.

A related significant result also due to [6] gives the point at which the alternative hypothesis is sufficiently removed from the null so that the distribution of $\lambda_1(\widehat{\Sigma})$ is stochastically different from that of $\lambda_1(\widehat{\Sigma})$ under the null. In that paper, under slightly different distributional assumptions (see Supplementary Material), the authors show that if the largest eigenvalue is changed from 1 to $\lambda_1(\Sigma) > 1 + \sqrt{p/n}$ and the other eigenvalues remain the same, then $\lambda_1(\widehat{\Sigma})$ has Gaussian fluctuations and they are of order $n^{-1/2}$. See also [41]. The value $1 + \sqrt{p/n}$ is therefore called the *Gaussian phase transition*. Part of our simulation study investigates whether the bootstrap is capable of capturing this statistically interesting phase transition.

Another very important result concerns the distribution of eigenvalues of the sample covariance matrix. If we call F_p the empirical spectral distribution of $\widehat{\Sigma}$, we have

Theorem 1.2 (Marchenko-Pastur, '67)**.** *Suppose $X_{i,j}$ are i.i.d with mean 0, variance 1, and a fourth moment. Then, as $p/n \to r \in (0,1)$,*

$$F_p \Longrightarrow F_r \ , a.s \ .$$

Furthermore, F_r, the so-called Marchenko-Pastur distribution, has the density f_r with

$$f_r(x) = \frac{\sqrt{(r_+ - x)(x - r_-)}}{2\pi r x} 1_{x \in (r_-, r_+)} \ ,$$

where $r_\pm = (1 \pm \sqrt{r})^2$.

Marchenko and Pastur's paper [36] contains many more results, including the case where the population covariance is not Id_p, the case where $r > 1$, etc. The previous result can be interpreted as saying that the histogram of sample eigenvalues is asymptotically non-random, and its limiting shape, which depends on the ratio p/n, is characterized by the Marchenko-Pastur distribution.

2 Simulation Study

We investigate via simulation the behavior of the bootstrap for the top eigenvalue of the standard sample covariance matrix, $\frac{1}{n-1}\tilde{X}'\tilde{X}$, where $\tilde{X}$ refers to the matrix of centered data values. For simplicity, we consider the case where only the top eigenvalue λ_1 is allowed to vary; we assume the remaining eigenvalues are all equal to 1. Therefore in this simple setting, the inferential question is to determine whether the top eigenvalue λ_1 differs from 1.

In our simulations, we generate data X for several distributions of X, and perform the bootstrap for the top eigenvalue of the sample covariance matrix. In what follows, we focus on when the observations X_i come from either a multivariate Normal distribution or an elliptical distribution with exponential weights (see Supplementary Text, Section S1 for details on this and other elliptical distributions we considered). Many theoretical results on random matrices in high dimensions do not yet extend to the case of elliptical distributions where the geometry of the data is more complex than under standard setups. Therefore, simulations under the elliptical distribution, while still an idealization, represents a simulation scenario that is somewhat more realistic than that of standard models.

Regarding inference on λ_1, we consider two different types of questions for which the bootstrap could be used. The first is to estimate basic features of the estimate $\hat{\lambda}_1$, such as its variance or bias. The second is to use the bootstrap to perform inference on λ_1, for example to create confidence intervals for λ_1; there exist several methods of creating bootstrap confidence intervals which we evaluate. The results of our simulations shows that the bootstrap performs quite badly for all of these tasks as the ratio p/n grows, unless the top eigenvalue λ_1 is very large and thus highly separated from the other eigenvalues. Our theoretical work explains these phenomena.

Estimating Bias As noted above in the background (see Subsection 1.2 and the Supplementary Material), the eigenvalues of the sample covariance are biased in high dimensions for estimating the true eigenvalues. Before considering the bootstrap estimates of bias, we first note the importance of this bias in understanding the behavior of $\hat{\lambda}_1$. The bias can be substantial unless λ_1 is quite large, and the bias is clearly evident even for low ratios of $r = p/n$ if λ_1 is close to 1. Furthermore, this bias is more pronounced for elliptical distributions than the normal distribution (which can be explained through the results of [20, 39]). For example, for the null setting of $\lambda_1 = 1$, with a ratio of p/n as low as 0.01, we see a bias in $\hat{\lambda}_1$, overestimating the true λ_1 by 17% for the normal distribution, and 49% for the elliptical distribution with exponential weights (Supplemental Table S1-S4). As the ratio of p/n grows, the bias increases, with $\hat{\lambda}_1$ overestimating λ_1 by 1.88 and 14.93 when $\lambda_1 = 1$ for the normal distribution and the elliptical distribution with exponential weight, respectively when $p/n = 0.5$. The bias declines as λ_1 grows and becomes more separated from the remaining eigenvalues, especially relative to the size of λ_1; Subsection 3.2 provides an explanation for this phenomenon. But the bias remains for large ratios of p/n even for λ_1 well beyond the Gaussian phase transition $(1 + \sqrt{p/n})$, especially for non-normal distributions; for $p/n = 0.5$, when λ_1 is as large as $1 + 11\sqrt{p/n} \approx 8.78$, the bias of $\hat{\lambda}_1$ is 8.08 when X follows an elliptical distribution with exponential weights, and 2.01 for an elliptical distribution with normal weights. Any use of $\hat{\lambda}_1$ as an estimate of λ_1 must grapple with the problem of such a highly non-consistent estimator, making bootstrap methods for estimating the bias highly relevant.

The standard bootstrap estimate of bias is given by, $\bar{\lambda}_1^* - \hat{\lambda}_1$, where $\bar{\lambda}_1^* = \frac{1}{B}\sum_b \hat{\lambda}_1^{*b}$ is the mean of the bootstrap estimates of λ_1. Unfortunately, we see in simulations that this bootstrap estimate of bias is not a reliable estimate of the bias of $\hat{\lambda}_1$ unless λ_1 is quite large relative to

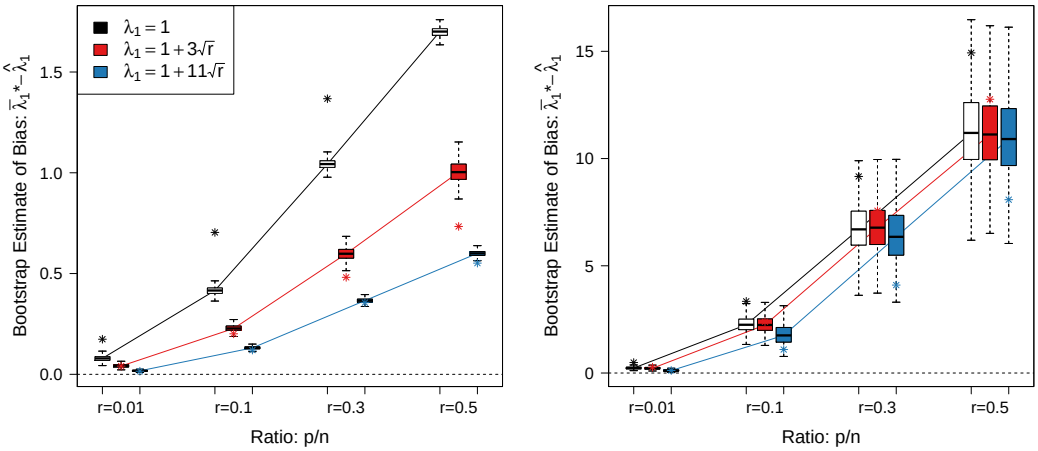

(a) $Z \sim$ Normal (b) $Z \sim$ Ellip. Exp

Figure 1: **Bias of Largest Bootstrap Eigenvalue, n=1,000:** Plotted are boxplots of the difference of $\bar{\lambda}_1^*$, the average bootstrap value of λ_1 over 999 bootstrap samples, minus the estimate $\hat{\lambda}_1$. This is repeated over 1000 simulations. $\bar{\lambda}_1^* - \hat{\lambda}_1$ is also the standard bootstrap estimate of bias. Each group of boxplots along the x-axis corresponds to a different ratio r of p/n; different colors of the boxplot correspond to different values of the true $\lambda_1 : 1, 1+3\sqrt{r}$,$1+11\sqrt{r}$; for larger values of λ_1 see Supplementary Figure S3. The asterisk (*) in the plot corresponds to the true bias, $\hat{\lambda}_1 - \lambda_1$ as evaluated over 1,000 simulations. See Supplemental Tables S1-S4 for the median values of these boxplots and for those of larger λ_1 values.

the other eigenvalues. We demonstrate these results in Figure 1 where we plot boxplots of the standard bootstrap estimates of bias over 1,000 simulations for different ratios of p/n and values of λ_1. For X_i following a normal distribution, the bootstrap estimate of bias remains poor for large p/n even for λ_1 past the phase transition, e.g. $\lambda_1 = 1 + 3\sqrt{p/n}$ ($\lambda_1 \approx 3.1$ for $p/n = 0.5$), and only for much larger values of λ_1 does the bootstrap estimate of bias start to approach the true bias in high dimensions (Supplementary Table S1). Further the bootstrap estimate of bias is inconsistent: depending on the true value of λ_1 the bootstrap either under- or over-estimates the bias. Another important feature of the bootstrap shown in our simulations is that when X_i follows an elliptical distribution with exponential weights, while the mean performance of the bootstrap estimate is still poor, it is the extremely high variance of the bootstrap that is even more problematic for the bootstrap estimate of bias. Indeed, as we discuss below, the bootstrap distribution of the top eigenvalue seems to change dramatically from simulations to simulations, thus creating highly variable estimates.

Estimating the Variance We see similar problems in our simulations for the boostrap estimate of variance (Figure 2). Specifically, the bootstrap dramatically overestimates the variance of $\hat{\lambda}_1$ when λ_1 is close to 1. When the X_i's are normally distributed and $\lambda_1 = 1$, the bootstrap estimates the variance to be four times larger than the true variance for $p/n = 0.1$, and grows to be up to 60 times larger than the true variance when $p/n = 0.5$ (Supplementary Table S5). Even when $\lambda_1 = 1 + 3\sqrt{p/n}$, well beyond the Gaussian phase transition and hence in a relatively easy setup,

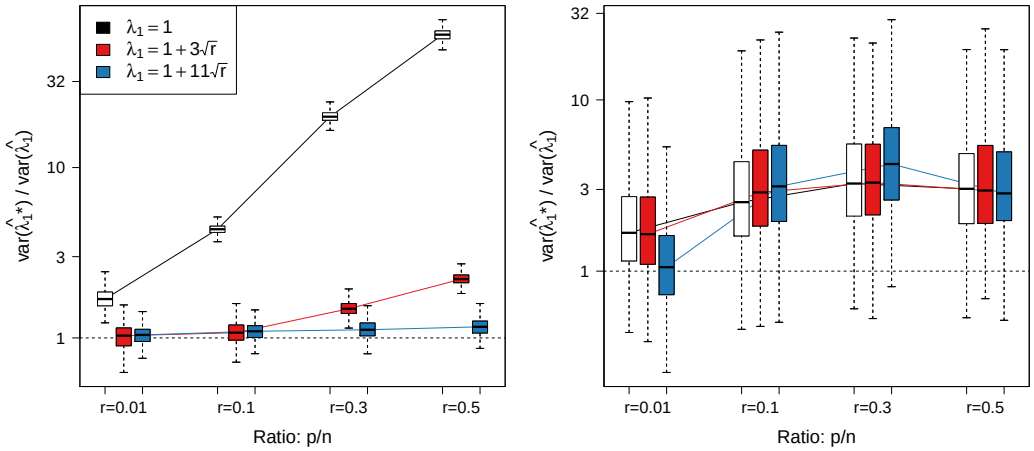

(a) $Z \sim$ Normal (b) $Z \sim$ Ellip. Exp

Figure 2: **Ratio of Bootstrap Estimate of Variance to True Variance for Largest Eigenvalue, n=1,000:** Plotted are boxplots of the bootstrap estimate of variance ($B = 999$) as a ratio of the true variance of $\hat{\lambda}_1$; boxplots represent the bootstrap estimate of variance over 1000 independent simulations. Each group of boxplots along the x-axis corresponds to a different ratio r of p/n; different colors of the boxplot correspond to different values of the true $\lambda_1 : 1$ (white), $1 + 3\sqrt{r}$ (red), and $1 + 11\sqrt{r}$ (blue); for larger values of λ_1 see Supplementary Figure S4. See Supplemental Table S5 for the median values of boxplots.

the bootstrap estimate of variance is inflated to 1.5 to 2.2 times as large as the truth, for $p/n = 0.3$ and 0.5, respectively. Only for large values of λ_1 does the bootstrap inflation of the variance become minimal. Again, when the X_i's follow an elliptical distribution with exponential weights, the behavior of the bootstrap estimate of variance is dominated by the variability in the estimate, because the distribution of $\hat{\lambda}_1^*$ is so erratic.

Confidence Intervals for λ_1 Standard techniques for creating confidence intervals for λ_1 are clearly problematic in high dimensions since $\hat{\lambda}_1$ is biased and not a consistent estimator for λ_1. If $\lambda_1(\Sigma)$ is beyond the Gaussian phase transition, it is nonetheless fairly straightforward to construct those intervals in the Wishart data setting. However, it is still useful to consider the performance of bootstrap confidence intervals because of what they highlight about the behavior of the bootstrap for the top eigenvalue. Another reason is that many practitioners might use the bootstrap to make statements about whether top eigenvalues are separated from the rest of the eigenvalue spectrum.

Bootstrap confidence intervals can be created in multiple ways [13]. Common techniques include 1) a simple normal confidence interval around $\hat{\lambda}_1$ using the bootstrap estimate of variance, 2) the percentile method, which uses the percentiles of the bootstrap distribution of $\hat{\lambda}_1^{*b}$, or 3) a bias-corrected confidence interval. Based on our earlier discussions, it is not surprising that none of these methods for estimating confidence intervals will have the proper coverage probability. Examining the actual bootstrap distributions of $\hat{\lambda}_1^* - \hat{\lambda}_1$ from multiple simulations (Figure 3), we see clearly the incorrect bias estimation and the overestimation of variance that results from using the bootstrap, features that will also invalidate confidence intervals constructed from the percentiles of

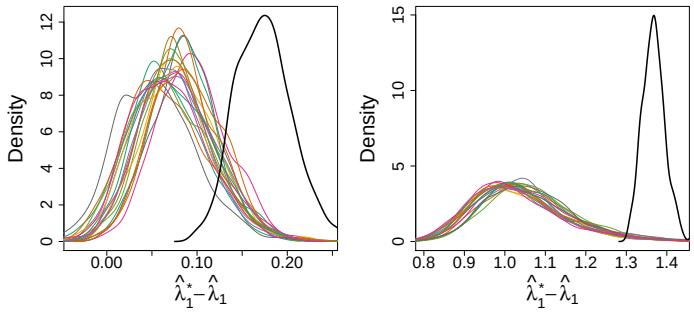

(a) $Z \sim$ Normal, r=0.01 (b) $Z \sim$ Normal, r=0.3

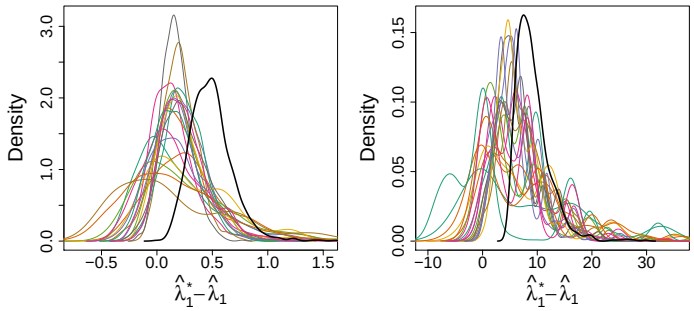

(c) $Z \sim$ Ellip. Exp, r=0.01 (d) $Z \sim$ Ellip. Exp, r=0.3

Figure 3: **Bootstrap distribution of $\hat{\lambda}_1^*$ under the null ($\lambda_1 = 1$), n=1,000:** Plotted are the estimated density of twenty simulations of the bootstrap distribution of $\hat{\lambda}_1^{*b} - \hat{\lambda}_1$, with $b = 1, \ldots, 999$. The solid black line line represents the distribution of $\hat{\lambda}_1 - \lambda_1$ over 1,000 simulations. For similar figures for the larger value of $\lambda_1 = 1 + 3\sqrt{r}$, see Supplementary Figure S6.

the distribution of $\hat{\lambda}_1^*$. We also see that when the X_i's follow an elliptical distribution with exponential weights, the bootstrap distributions do not appear to be converging to a limit for small values of λ_1, at least for $n = 1000$ – an even greater problem in using the bootstrap in these settings.

As expected, the resulting bootstrap confidence intervals are not useful in inference on the true value of λ_1. Bootstrap confidence intervals based on the percentile estimates do not cover the true value with any kind of reasonable probability until λ_1 becomes quite large (Figure 4 and Supplementary Table S6); this is undoubtedly because of the bias in the estimate of λ_1 making the percentile method inappropriate for constructing confidence intervals. Bootstrap confidence intervals based on normal intervals around $\hat{\lambda}_1$ using the bootstrap estimate of variance *do* cover the true value of the λ_1 with high probability. However, this coverage is due to the fact that the bootstrap estimate of variance is much larger than the true variance of $\hat{\lambda}_1$, as seen above, and thus result in overly large confidence intervals. As a result, the normal-based bootstrap confidence intervals also incorrectly cover the putative null hypothesis ($\lambda_1 = 1$) with high probability when the alternative is true, particularly for X_i following an elliptical distribution (Supplementary Table S7). In short, such bootstrap confidence intervals based on the bootstrap estimate of variance suffer from lack of power once the distribution of X_i deviates from strictly normal because of the large size of the confidence interval.

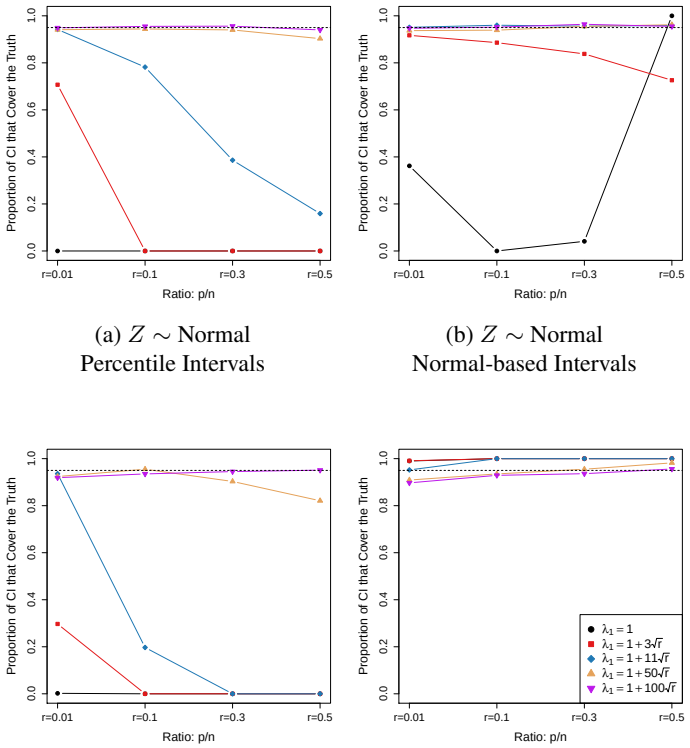

(a) $Z \sim$ Normal
Percentile Intervals

(b) $Z \sim$ Normal
Normal-based Intervals

(c) $Z \sim$ Ellip. Exp
Percentile Intervals

(d) $Z \sim$ Ellip. Exp
Normal-based Intervals

Figure 4: **95% CI Coverage,** $n = 1,000$**:** Plotted are the percentage of the confidence intervals that cover the true λ_1 (out of 1,000 simulations), for different values of $r = p/n$ and for different true values of λ_1. The plotted values can be found in Supplementary Table S6, and the percentage of these same intervals that cover the null value $\lambda_1 = 1$ can be found in Supplementary Table S7. Additional elliptical distributions can be seen in S5.

2.1 Statistics for Detecting Gaps in the Eigenvalue Spectrum

The behavior of the top eigenvalues is often studied in theoretical work, but in practice, examination of the eigenvalues of the sample covariance matrix is largely done to find gaps in the eigenvalue spectrum. Such gaps might indicate a logical point at which to reduce the dimension of the data. Again, we focus for simplicity on detecting the separation of just the top eigenvalue from the remainder. Then a natural statistic is the gap statistic, $\hat{\lambda}_1 - \hat{\lambda}_2$, where large values of the statistic are meant to suggest that there is a large difference between the first and second eigenvalue.

These statistics are difficult to understand theoretically, with limit distributions that are even less standard than the Tracy-Widom distribution (see [47], which explains joint distributional results, [15], and [20] for applications), which again makes them good candidates for using the bootstrap for inference. We consider the performance of the bootstrap for these statistics via the same simulation structure that we applied to the largest eigenvalue, above.

The gap statistic $\hat{\lambda}_1 - \hat{\lambda}_2$ is also a biased estimate for the true population value. And unlike $\hat{\lambda}_1$, the direction of the bias for the gap statistic differs depending on the value of λ_1. For $\lambda_1 = 1$, the

gap statistic overestimates the true difference, while for $\lambda_1 > 1$, the gap statistic underestimates the true difference (Supplementary Tables S8-S11); how large λ_1 needs to be before the bias becomes negative depends on the distribution of X.

As in the case of the top eigenvalue, the bootstrap estimate of bias does not accurately estimate this bias (Supplementary Figure S7). The bootstrap under-estimates the absolute size of the bias, and for elliptical distributions can misspecify the direction of the bias (Supplementary Tables S8-S11, Supplementary Figure S7). As with the top eigenvalue, the disparity in the bootstrap estimate of bias improves as the top eigenvalue becomes more separated from the bulk. Estimating the variance of the gap statistic with the bootstrap shows similar problems, with the bootstrap widely over-estimating the variance of $(\lambda_1(\widehat{\Sigma}) - \lambda_2(\widehat{\Sigma}))$ in high dimensions (Supplementary Figure S8). Bootstrap confidence intervals also suffer from the same problem as those of the top eigenvalue: percentile CIs have low coverage of the truth in high dimensions and normal-based CI being much wider than necessary because of the over estimation of the variance.

Gap Ratio Statistic Another alternative that tries to normalize the gap statistic is the gap ratio, $(\lambda_1 - \lambda_2)/(\lambda_2 - \lambda_3)$. Onatski [38] proposed tests based on these statistics to avoid having to estimate $\sigma_{n,p}$ and $\mu_{n,p}$ in Theorem 1.1, when Σ is a multiple of the identity. In the scenario we are evaluating, this population quantity is not well defined ($\lambda_2 - \lambda_3 = 0$), but the estimate and its distribution are well defined, and the statistic is a tool for deciding whether λ_1 and λ_2 are well separated. Again, the bootstrap estimate gives poor estimates of various features of the actual distribution of the gap ratio statistic: the bootstrap estimate is biased and can either under or over estimate the variance, depending on the true value of λ_1 (Supplementary Figures S11 and S12). The bootstrap distribution does not appear to be converging, i.e the bootstrap distribution seems to change with each X (Supplementary Figure S13).

3 Theoretical results

The problems with the bootstrap can be explained in part by the difference between the spectral behavior of weighted and unweighted covariance matrices when p/n is not small. Specifically, bootstrapping the observations (rows) of X is equivalent to randomly reweighting the observations which changes the spectral distribution of $\widehat{\Sigma}$. For example, if X_i's are normally distributed, randomly weighting the X_i's transforms the data to an elliptical distribution, which leads to a very different spectral distribution of eigenvalues when p/n is not close to 0 (see Theorem S2.2 in the Supplementary material). Similarly, the distribution of the largest eigenvalues are dramatically affected by reweighting; the one exception to this rule is the situation where the largest eigenvalues of Σ are very separated from the rest.

In what follows we provide theoretical results that help explain the results of our numerical simulations and also complete them. The first set of results concerns the impact of bootstrapping on the spectral distribution of a sample covariance matrix. We explain that this creates bias in the setting we consider and it helps explain some of the misbehavior of the bootstrap we observed in the numerical study. We then consider the case of extreme eigenvalues, in the case where the largest population eigenvalues are well-separated from the bulk of the eigenvalues. We show that then the bootstrap works asymptotically under certain conditions. This helps explain why the performance of the bootstrap improves in our numerical study when we increase the largest population eigenvalue.

3.1 Bootstrapped empirical distribution

3.1.1 A theoretical result

Lemma 3.1. *Suppose $\{X_i\}_{i=1}^n$ are fixed vectors in $\mathbb{R}^p$. Suppose w_i's are independent random variables. Consider $S_w = \frac{1}{n}\sum_{i=1}^n w_i X_i X_i'$ and call*

$$m(z) = \frac{1}{p}\textit{trace}\left((S_w - z\mathrm{Id}_p)^{-1}\right), \textit{ for } z = u + iv \in \mathbb{C}^+ .$$

Then

$$P(|m(z)) - \mathbf{E}\left(m(z)\right)| > t) \leq C\exp(-cp^2v^2t^2/n) ,$$

with $C = 4$ and $c = 1/16$ for instance.

The same result holds when w_i's have $Mult(n, 1/n)$ distribution.

Naturally, the case of $Mult(n, 1/n)$ corresponds to the standard bootstrap. We explore the statistical implications of this result in Section 3.1.2.

Corollary 3.1. *When $p/n \to \kappa \in (0, \infty)$, the Stieltjes transform $m(z)$ of the independently-weighted bootstrapped covariance matrix is asymptotically deterministic. The same is true with the standard bootstrap, where the weights have a Multinomial(n,1/n) distribution.*

In particular, if f is a bounded continuous function, as n and p tend to infinity, while $p/n \to \kappa$,

$$\frac{1}{p}\sum_{i=1}^p f(\lambda_i^*) - \frac{1}{p}\sum_{i=1}^p \mathbf{E}^*\left(f(\lambda_i^*)\right) \to 0 \textit{ in probability} ,$$

where λ_i^* are the decreasingly ordered bootstrapped eigenvalues and $\mathbf{E}^*(\cdot)$ refers to expectation under the bootstrap distribution.

The corollary follows from our lemma simply by using the well-known fact that convergence of the Stieltjes transform implies convergence of the corresponding spectral distributions (see [5], [25]).

3.1.2 Statistical consequences when p/n is not small but remains bounded

One of the main problem that the bootstrap exhibits in the setting we consider is bias. Let us give a concrete example.

Bias in bootstrap spectral distribution: the case of Gaussian data Suppose $X_i \overset{iid}{\backsim} \mathcal{N}(0, \Sigma)$ and suppose that $\widehat{\Sigma} = \frac{1}{n}\sum_{i=1}^n X_i X_i'$. When bootstrapping, we effectively observe $\widehat{\Sigma}_w = \frac{1}{n}\sum_{i=1}^n w_i X_i X_i'$, where w_i is the number of times index i is picked in our resample. It is known that the spectral distribution of $\widehat{\Sigma}$, $L_n(\widehat{\Sigma})$, has a non-random limit, $\mathcal{L}(\Sigma)$ satisfying the so-called Marchenko-Pastur equation (see [36],[50], [46],[5] and Theorem S2.1 in the Supplementary Material). It is also known that $\widehat{\Sigma}_w$ when sampling both w_i's and X_i's has a non-random limiting distribution (for instance when $\{w_i\}_{i=1}^n$ are independent of $\{X_i\}_{i=1}^n$), $\mathcal{L}(\Sigma, w)$, which can be implicitly characterized by a pair of equations for a variant of the Stieltjes transform of $\widehat{\Sigma}_w$ (see [11], [21] and Theorem S2.2 in the Supplementary Material). This limit distribution is in general hard to characterize analytically,

however it is not the same as that of $\widehat{\Sigma}$, i.e $\mathcal{L}(\Sigma, w) \neq \mathcal{L}(\Sigma)$. Furthermore, it is clear by a simple conditioning argument, that, if $L_n(\widehat{\Sigma}_w)$ is the spectral distribution of $\widehat{\Sigma}_w$,

$$L_n(\widehat{\Sigma}_w)|\{X_i\}_{i=1}^n \Longrightarrow \mathcal{L}(\Sigma, w)a.s$$

where the $a.s$ statement refers to the design matrix.

$L_n(\widehat{\Sigma}_w)|\{X_i\}_{i=1}^n$ is the bootstrapped spectral distribution of $\widehat{\Sigma}$. Its limit, $\mathcal{L}(\Sigma, w)$ is different from that of $\widehat{\Sigma}$, $\mathcal{L}(\Sigma)$. *Hence the bootstrapped distribution of the eigenvalues of $\widehat{\Sigma}$ is in general biased.*

In connection with the results of [42], these results also help explain the problem of bias found in the bootstrapped extreme eigenvalues, when for instance $\Sigma = \mathrm{Id}_p$, i.e no eigenvalues are well-separated from the bulk.

Extensions Much work has been done in random matrix theory to extend the domain of validity of the Marchenko-Pastur equation, which holds beyond the case of Gaussian data. The bootstrap bias problem remains the same, because the limiting properties of the matrices of interest are unaffected by the move from Gaussian to these more general models (see [21]).

Bootstrap and geometry: an explanation of the problems with the bootstrap At a high-level, one can intuitively think that bootstrapping moves the data in the setting considered here from a Gaussian setting to an elliptical one. It is well-known ([14], [29], [21]) that in moderate and high-dimension elliptical distributions have completely different geometric properties than Gaussian ones and that this geometric features impact strongly the statistical behavior of many estimators, and in particular that of eigenvalues of sample covariance matrices ([21]). As such, it is not that surprising that the bootstrap does not perform well: from an eigenvalue point of view, it is as if the bootstrap changed "the geometry of the dataset" and this geometry has an important impact on their behavior.

Bootstrapping is therefore not a good way to mimick the data generating process in this context.

Other designs Lemma 3.1 and Corollary 3.1 apply without restrictions on the design, which is one of their main strength. With further specifications, for instance assuming that the data is generated from an elliptical distribution, we could characterize precisely various aspects of the bootstrapped spectral distribution in relation to the empirical spectral distribution, using for instance Theorem S2.2. However, this would be quite model-specific, and since we are focused on understanding more generally the problems with the bootstrap, we do not think these simple computations would serve our purpose. So we do not detail them here.

3.2 Extreme eigenvalues

Definition 1. Suppose $\widehat{\theta}_n(X_1, \ldots, X_n)$ is a statistic, $\widehat{\theta}_n^*$ is its bootstrapped version. Suppose that $\widehat{\theta}_n \Longrightarrow T$. We say that *the bootstrap is consistent in probability* if

$$\widehat{\theta}_n^* \Longrightarrow_w T \quad \text{in } P_{X_1,\ldots,X_n} - \text{probability}.$$

Here $\Longrightarrow_w$ refers to weak convergence of $\widehat{\theta}_n^*$ under the bootstrap weight distribution; the convergence in probability is with respect to the joint distribution of $X_1, \ldots, X_n$, which we denote $P_{X_1,\ldots,X_n}$. For simplicity, we often abbreviate $P_{X_1,\ldots,X_n}$ by P.

We say that *the bootstrap is strongly consistent* if

$$\widehat{\theta}_n^* \Longrightarrow_w T \quad \text{a.s } P_{X_1,\ldots,X_n}.$$

3.2.1 Approximation results

Call S_n the sample covariance matrix of the data. We use the block notation

$$S_n = \begin{pmatrix} T_n & U_n \\ U_n' & n^{-\alpha}V_n \end{pmatrix}$$

Call Σ_n the true covariance. We use the block notation

$$\Sigma_n = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & n^{-\alpha}\Sigma_{22} \end{pmatrix} .$$

T_n and Σ_{11} are both assumed to be $q \times q$.

Assumptions

- **A1** We assume that $|||\Sigma_{22}|||_2 = \mathrm{O}(1)$ and that $\lambda_{min}(\Sigma_{11}) > \eta > 0$. We assume that Σ_{11} is $q \times q$ with q fixed.
- **A2** X_i's are i.i.d with $X_i = r_i Z_i$, where $Z_i \sim \mathcal{N}(0, \Sigma_n)$, and $0 < \delta_0 < r_i < \gamma_0$ is a bounded random variable independent of Z_i, with $\mathbf{E}\left(r_i^2\right) = 1$.
- **A3** The bootstrap weights w_i have infinitely many moments, $\|w\|_\infty = \mathrm{O}(\mathrm{polyLog}(n))$ and $\mathbf{E}\left(w_i\right) = 1$. These weights can either be independent or Multinomial$(n, 1/n)$.
- **A4** p/n remains bounded as n and p tend to infinity.

We then have the following theorem.

Theorem 3.1. *Under our assumptions A1-A4, if $\alpha > 1/2 + \epsilon$ for any $\epsilon > 0$,*

$$\sup_{1 \le i \le q} \sqrt{n}(\lambda_i(S_n) - \lambda_i(T_n)) = \mathrm{o}_P(1) . \tag{1}$$

Furthermore, if w denotes the vector of weights used in the bootstrap and the corresponding bootstrapped matrices are S_n^* and T_n^*, we have

$$\sup_{1 \le i \le q} \sqrt{n}(\lambda_i(S_n^*) - \lambda_i(T_n^*)) = \mathrm{o}_{P,w}(1) . \tag{2}$$

3.2.2 Consequences for the bootstrap

Recall some of the key results of [7, 8] and [17], p.269: in the low-dimensional case, where p is fixed and $n \to \infty$, if all eigenvalues of Σ are simple, the bootstrap distribution of the eigenvalues of S_n is strongly consistent. On the other hand, it is known from these papers that the bootstrap distribution of the eigenvalues of S_n is inconsistent when the eigenvalues of Σ have multiplicities higher than 1.

In light of these results, we have the following theorem

Theorem 3.2. *Suppose the eigenvalues of Σ_{11} are simple and the assumptions A1-A4 of Theorem 3.1 hold. Then the bootstrap distribution of the q largest eigenvalues of S_n is consistent in probability.*

3.2.3 Discussion of assumptions and remarks

$\alpha > 1/2 + \epsilon$ This assumption is not terribly restrictive: in fact under our assumptions, if $\Sigma_{22} = \mathrm{Id}_{p-q}$, the fraction of variance explained by the top q eigenvalues is, if $C = \text{trace}\,(\Sigma_{11})/q$

$$\frac{\text{trace}\,(\Sigma_{11})}{\text{trace}\,(\Sigma_n)} = \frac{qC}{qC + (p-q)n^{-\alpha}}\,.$$

So if $1/2 + \epsilon < \alpha < 1$, in our asymptotics where p/n remains bounded, this fraction of variance is asymptotically 0, since $pn^{-\alpha} \to \infty$ and qC is bounded. On the other hand, if $\alpha > 1$, the fraction of variance explained by the top q eigenvalues is approximately 1, which is the standard setting for the use of PCA.

If $\alpha = 1$, the fraction of variance varies between 0 and 1, depending on C. Of course, a similar analysis is possible and actually easy to carry out if Σ_{22} is not a multiple of the identity. We leave those details to the interested reader.

Strong consistency of the bootstrap We have chosen to present our results using convergence in probability statements, as we think they better reflect the questions encountered in practice. However, a quick scan through the proofs show that all the approximation results could be extended to a.s convergence: the random matrix results we rely on hold a.s, and the low-dimensional bootstrap results we use also hold a.s in low-dimension.

Distributional assumptions on X_i's The assumption that $Z_i \sim \mathcal{N}(0, \Sigma)$ is not critical: most of our arguments could be adapted to handle the case where $Z_i = \Sigma^{1/2} Y_i$, where Y_i has independent, mean 0, variance 1 entries, with sufficiently many moments. This simply requires appealing to slightly different random matrix results that exist in the literature. Also, the first q coordinates of X_i could have a much more general distribution than the elliptical distributions we consider here, as our proof simply requires control of $|||V_n|||_2$, which is where we appeal to random matrix theory. Doing this entails minor technical modifications to the proof, but since it might reduce clarity, we leave them to the interested reader.

Assumptions on Σ The block representation assumptions are made for analytic convenience and can be easily dispensed of: eigenvalues are of course unaffected by rotations, so we simply chose to write Σ in a basis that gave us this nice block format. As long as the ratio between the q-th eigenvalue of Σ and $q+1$st is of order n^α, our results hold. Furthermore, our results also handle a situation similar to ours, where for instance the top q largest eigenvalues of Σ grow like n^α and the $q+1$st is of order 1, by simple rescaling, using for instance the fact that $\text{trace}\,(T_n)/\text{trace}\,(\Sigma_{11}) \to 1$ in probability.

The m-out-of-n bootstrap As noted in [8, 17], subsampling approaches fix the problem of bootstrap inconsistency in the setting where Σ_{11} has eigenvalues of multiplicity higher than 1. Our approximation results for the pair (S_n^*, T_n^*) can be extended to subsampling approaches, and hence our results could be extended to cover these ideas. However, since this question is a bit distant from our main motivations, we do not treat it in detail.

3.2.4 Eigenvalues that are not well separated

Our results on the bootstrapped empirical spectral distribution apply here and in particular suggests that problems are likely to arise in practice. However, interesting questions concerning the fluctuation behavior of bootstrapped eigenvalues when the largest population eigenvalues are not

well-separated from the bulk are very natural. For instance, is it the case that the bootstrap distribution of the largest eigenvalue of a sample covariance matrix is Tracy-Widom when that is the case for the sampling distribution of the top eigenvalue?

Though mathematically interesting, this question does not seem so statistically important in light of the simulation study in Section 2. For instance Figure 1 indicates that the bootstrap estimate of bias of $\lambda_1(\widehat{\Sigma})$ is itself very biased. Figure 3 and in particular Subfigure 1b suggests that the bootstrap distribution of our statistics is a poor approximation of the sampling distribution. It also suggests that the characteristics of the bootstrap distribution may depend strongly on the property of the design matrix, rendering their characterization mathematically intractable outside of simple situations of limited practical statistical interest. For this reason, we postpone this mathematically interesting and delicate question to possible future work.

4 Conclusion

We have investigated in this paper the properties of the bootstrap for spectral analysis of high-dimensional covariance matrices. We have shown numerically and through theoretical results that in general, the bootstrap does not provide accurate inferential results. The one exception of practical interest is the situation where we are interested in inference concerning only a few large eigenvalues of Σ, which are well-separated from the bulk of them and of multiplicity 1. In this case, the problem is effectively low dimensional (Theorem 3.1), and the bootstrap works, because it is known to work in low-dimension when Σ has eigenvalues of multiplicity 1.

This confirms the findings of [22]: in high-dimension, bootstrapping does not mimick the data generating process. Hence, standard bootstraps appear to work only for problems that are effectively low-dimensional.

SUPPLEMENTARY MATERIAL

Supplementary Text More detailed description of simulations and proofs of the theorems stated in main text (see below; see also authors' website for different formatting)

Supplementary Figures Supplementary Figures referenced in the main text (pdf; see also authors' website for different formatting)

Supplementary Tables Supplementary Tables referenced in the main text (pdf; see also authors' website for different formatting)

APPENDIX

S1 Description of Simulations and other Numerics

For each of 1,000 simulations, we generate a $n \times p$ data matrix X. For each X, we calculate either the top eigenvalue (or the gap statistic) from the sample covariance matrix. Specifically, we perform the SVD of X using the ARPACK numeric routines (implemented in the package `rARPACK` in R) to find the top five singular values of X and get estimates $\hat{\lambda}_i$ by multiplying the singular values of X by $1/n$.

For each simulation, we perform bootstrap resampling of the n rows of X to get a bootstrap resample X^{*b} and λ_i^{*b}; we repeat the bootstrap resampling $B = 999$ times for each simulation, resulting in 999 values of λ_i^{*b} for each simulation.

In all simulations, we only consider the case where only λ_1 is allowed to differ from the rest of the eigenvalues. Therefore for all eigen values except the first, $\lambda_i = 1$. For λ_1, we consider $\lambda_1 = 1 + c\sqrt{\frac{p}{n}}$ for the following c: 0, 0.9, 1.1, 1.5, 2.0, 3.0, 6.0, 11.0, 50.0, 100.0, and 1000.0 (not all of these values are shown in figures or tables accompanying this manuscript). The results shown in this manuscript set $n = 1,000$, though $n = 500$ was also simulated.

Generating X We generate X as $X = Z\Sigma$ where $\Sigma = V\Lambda V'$, and Λ is a diagonal matrix of eigenvalues $\lambda_1 \geq \ldots \geq \lambda_p$. We assume that there is no structure in the true eigenvectors, and generate V as the right eigenvectors of the SVD of a nxp matrix with entries i.i.d $N(0,1)$.

$Z = DZ_0$ is a nxp matrix, with Z_0 having entries i.i.d. $N(0,1)$ and D is a diagonal matrix D. If D is the identity matrix, Z will be i.i.d. normally distributed; otherwise Z will be i.i.d with an elliptical distribution. We simulated under the following distributions for the diagonal entries of D to create elliptical distribution for Z,

- $D_{ii} \sim N(0,1)$
- $D_{ii} \sim Unif(1/2, \frac{\sqrt{3}\sqrt{4-1/4}}{2} - \frac{1}{4})$
- $D_{ii} \sim Exp(\sqrt{2})$

In this manuscript, we concentrated only on $D_{ii} \sim Exp(\sqrt{2})$, the "Elliptical Exponential" distribution. This was because its behavior resulted in an elliptical distribution for Z with properties the most different from when Z is normal. The remaining choices for the distribution of D result in elliptical distributions between that of the Elliptical Exponential and the Normal. The results from when $D_{ii} \sim Unif$ were generally fairly similar to when Z is normal and results from when $D_{ii} \sim N(0,1)$ were more different, though not as extreme as the exponential weights.

S2 Review of existing results in random matrix theory

Notations We call $\lambda_1(M)$ or $\lambda_{max}(M)$ the largest eigenvalue of a symmetric matrix M. We call $\lambda_1(M) \geq \lambda_2(M) \geq \lambda_3(M) \geq \lambda_p(M)$ the ordered eigenvalues of the $p \times p$ matrix M. If

$Z \sim \mathcal{N}_{\mathbb{C}}(0, \Sigma)$, Z has a complex normal distribution, i.e $Z = \frac{1}{\sqrt{2}}(Z_1 + iZ_2)$ where Z_1 and Z_2 are independent with $Z_i \sim \mathcal{N}(0, \Sigma)$. We call $\mathbb{C}^+$ the set of complex numbers with positive imaginary part.

S2.1 Bulk results

Bulk results are concerned with the spectral distribution of $\widehat{\Sigma}$, i.e the (random) probability measure with distribution

$$dF_p(x) = \frac{1}{p}\sum_{i=1}^{n} \delta_{\lambda_i(\widehat{\Sigma})}(x) \ .$$

An efficient way to characterize the limiting behavior of F_p is through its Stieltjes transform:

$$\text{for } z = u + iv \text{ with } v > 0, \quad m_p(z) = \frac{1}{p}\text{trace}\left((\widehat{\Sigma} - z\text{Id}_p)^{-1}\right) \ .$$

Note that $m_p(z) : \mathbb{C}^+ \mapsto \mathbb{C}^+$. We have of course

$$m_p(z) = \int \frac{dF_p(\lambda)}{\lambda - z} = \frac{1}{p}\sum_{i=1}^{p} \frac{1}{\lambda_i(\widehat{\Sigma}) - z} \ .$$

An important result in this area is the so-called Marchenko-Pastur equation [36, 46], which states the following

Theorem S2.1. *Suppose $X_i \overset{iid}{\backsim} \Sigma^{1/2} Z_i$, where Z_i has $i.i.d$ entries, with mean 0, variance 1 and 4 moments. Suppose that the spectral distribution of Σ has a limit H in the sense of weak convergence of probability measures and $p/n \to r \in (0, 1)$. Then*

$$F_p \Longrightarrow F a.s \ ,$$

where F is a deterministic probability distribution.

Call $v_p(z) = (1 - p/n)\frac{-1}{z} + \frac{p}{n}m_p(z)$. Then $v_p(z) \to v_F(z)$ a.s. The Stieltjes transform of F can be characterized through the equation

$$-\frac{1}{v_F(z)} = z - r\int \frac{\lambda dH_\infty(\lambda)}{1 + \lambda v_F(z)} \ , \forall z \in \mathbb{C}^+$$

At an intuitive level, this result means that the histogram of eigenvalues of $\widehat{\Sigma}$ is asymptotically non-random. Its shape is characterized by F and its Stieltjes transform, m.

A generalization of this result to the case of elliptical predictors was obtained in [21]. For the purpose of the current paper, the main result of [21] states the following:

Theorem S2.2. *Suppose $X_i \overset{iid}{\backsim} \Sigma^{1/2} Z_i$, where Z_i has $i.i.d$ entries, with mean 0, variance 1 and 4 moments. Suppose that the spectral distribution of Σ has a limit H, that H has one moment and $p/n \to r \in (0, \infty)$. Consider the matrix*

$$B_n = \frac{1}{n}\sum_{i=1}^{n} w_i X_i X_i' \ .$$

Assume that the weights $\{w_i\}_{i=1}^n$ are independent of X_i's. Call ν_n the empirical distribution of the weights w_i's and suppose that $\nu_n \Longrightarrow \nu$.

Then $B_n \Longrightarrow B$ a.s, where B is a deterministic probability distribution; furthermore the Stieltjes transform of B, m, satisfies the system

$$m(z) = \int \frac{dH(\tau)}{\tau \int \frac{w^2}{1+rw^2\gamma(z)} d\nu(w) - z} \text{ and}$$
$$\gamma(z) = \int \frac{\tau dH(\tau)}{\tau \int \frac{w^2}{1+rw^2\gamma(z)} d\nu(w) - z} .$$

where $\gamma(z)$ is the only solution of this equation mapping $\mathbb{C}^+$ into $\mathbb{C}^+$.

Theorem S2.2 is interesting statistically because it shows that the limiting spectral distribution of weighted covariance matrices is completely different from that of unweighted covariance matrices, even when $\mathbf{E}(w_i) = 1$. This is in very sharp contrast with the low-dimensional case. (Note that as shown in [20], Theorem S2.2 holds for many other distributions for X_i's that the one mentioned in our statement.)

In the context of the current paper, this result is especially useful since bootstrapping a covariance matrix amounts to moving from an unweighted to a weighted covariance matrix.

S2.2 Edge results

Edge results are concerned with the fluctuation behavior of the eigenvalues that are at the edge of the spectrum of the matrices of interest.

A now standard result is due to Johnstone [32].

Theorem S2.3. *Suppose $X_i \overset{iid}{\backsim} \mathcal{N}(0, \mathrm{Id}_p)$. Assume that $p/n \to r \in (0,1)$. Then*

$$n^{2/3} \frac{\lambda_{max}(\widehat{\Sigma}) - \mu_{n,p}}{\sigma_{n,p}} \Longrightarrow TW_1 .$$

We have for instance $\mu_{n,p} = (1+\sqrt{p/n})^2$ and $\sigma_{n,p} = (1+\sqrt{p/n})(1+\sqrt{n/p})^{1/3}$. The result for the case $r \in (1,\infty)$ follows immediately by changing the role of p and n; see [32] for details.

Following a question of Johnstone, Baik, Ben-Arous and Péché obtained the following result [6].

Theorem S2.4. *Suppose $X_i \overset{iid}{\backsim} \mathcal{N}_{\mathbb{C}}(0, \Sigma)$. Suppose that $\lambda_1(\Sigma) = 1 + \eta\sqrt{p/n}$ and $\lambda_i(\Sigma) = 1$ for $i > 1$.*

1. *If $0 < \eta < 1$, then*
$$n^{2/3} \frac{\lambda_1(\widehat{\Sigma}) - \mu_{n,p}}{\sigma_{n,p}} \Longrightarrow TW_2 .$$

2. *On the other hand, if $\eta > 1$, then*
$$\sqrt{n} \frac{\lambda_1(\widehat{\Sigma}) - \mu_{\eta,n,p}}{\sigma_{\eta,n,p}} \Longrightarrow \mathcal{N}(0,1) .$$

Here,

$$\mu_{\eta,n,p} = \lambda_1(1 + \frac{\sqrt{p/n}}{\eta}) \text{ and } \sigma_{\eta,n,p} = \lambda_1\sqrt{1-\eta^{-2}} \ .$$

We note that we can rewrite the previous quantities solely as functions of λ_1, specifically

$$\mu_{\eta,n,p} = \lambda_1\left(1 + \frac{p/n}{\lambda_1 - 1}\right) \text{ and } \sigma_{\eta,n,p} = \lambda_1\sqrt{1 - \frac{n}{p}(\lambda_1-1)^{-2}} \ .$$

This representation shows that $\mu_{\eta,n,p}$ is an increasing function of λ_1 on $(1+\sqrt{p/n}, \infty)$ and therefore it would be easy to estimate $\lambda_1(\Sigma)$ from $\lambda_1(\widehat{\Sigma})$. In particular, it is very simple to build confidence intervals in this context.

Interpretation of Theorem S2.4 In other words, there is a phase-transition: if λ_1 is sufficiently large, i.e larger than $1+\sqrt{p/n}$ the largest eigenvalue of $\widehat{\Sigma}$ has Gaussian fluctuations. If it is not large enough, i.e smaller than $1+\sqrt{p/n}$, the fluctuations are Tracy-Widom, and in fact are the same if $\lambda_1(\Sigma) = 1$. Statistically, it is hard to build confidence intervals for λ_1 in the latter case - but it is very easy to do so in the first case where λ_1 is sufficiently large.

Similar results were obtained in [20] for general Σ in the complex Gaussian case and extended to the real case in [35]. [19] showed that Theorem S2.3 holds when $p/n \to 0$ and $p \to \infty$ at any rate. See also the interesting [41] and [31]. We finally note that the main result in [6] is slightly more general than Theorem S2.4 but we just need that version for the current paper.

S3 Proofs

S3.1 Proof of Lemma 3.1

Proof of Lemma 3.1. We recall the following result from a simple application of the Sherman-Morrison-Woodbury formula (see [30] and [5]): if M is a symmetric matrix, q is a real vector $v > 0$ and and $z = u + iv \in \mathbb{C}^+$,

$$|\text{trace}\left((M + qq' - z\text{Id}_p)^{-1}\right) - \text{trace}\left((M - z\text{Id}_p)^{-1}\right)| \leq \frac{1}{v} \ .$$

We use bounded martingale difference arguments as in [26], [40], [21].

• **Case 1: independent weights** w_i

Consider the filtration $\{\mathcal{F}_i\}_{i=0}^n$, with $\mathcal{F}_i = \sigma(w_1, \ldots, w_i)$ - the σ-field generated by $w_1, \ldots, w_i$ - and $\mathcal{F}_0 = \emptyset$.

Call $S_w^{(i)} = S_w - \frac{1}{n}w_i X_i X_i' - z\text{Id}_p$. In light of the result we just mentioned,

$$\frac{1}{p}|\text{trace}\left([S_w^{(i)}]^{-1}\right) - \text{trace}\left([S_w]^{-1}\right)| \leq \frac{1}{pv} \ .$$

In particular, this implies, since $\mathbf{E}\left(\text{trace}\left([S_w^{(i)}]^{-1}\right)|\mathcal{F}_i\right) = \mathbf{E}\left(\text{trace}\left([S_w^{(i)}]^{-1}\right)|\mathcal{F}_{i-1}\right)$ that

$$\frac{1}{p}|\mathbf{E}\left(\text{trace}\left([S_w]^{-1}\right)|\mathcal{F}_i\right) - \mathbf{E}\left(\text{trace}\left([S_w]^{-1}\right)|\mathcal{F}_{i-1}\right)| \leq \frac{2}{pv} \ .$$

Hence, $d_i = \frac{1}{p}\mathbf{E}\left(\text{trace}\left([S_w]^{-1}\right)|\mathcal{F}_i\right) - \frac{1}{p}\mathbf{E}\left(\text{trace}\left([S_w]^{-1}\right)|\mathcal{F}_{i-1}\right)$ is a bounded martingale-difference sequence. We can therefore apply Azuma's inequality ([34], p. 68), to get

$$P(|m_p(z) - \mathbf{E}\left(m_p(z)\right)| > t) \leq C\exp(-c\frac{p^2v^2t^2}{n}) \ .$$

In [21] it is shown that we can take $C = 4$ and $c = 1/16$.

• **Case 2: multinomial weights**

In this case, the previous result cannot be applied directly because the weights are not independent, since they must sum to n. However, to draw according to a Multinomial$(n, 1/n)$, we can simply pick an index from $\{1, \ldots, n\}$ uniformly and repeat the operation n times independently. Let $I(k)$ be the value of the index picked on the k-th draw from our sampling scheme. Clearly, the bootstrapped covariance matrix can be written as

$$S_w = \frac{1}{n}\sum_{k=1}^{n} X_{I(k)}X'_{I(k)} \ .$$

Consider the filtration $\{\mathcal{F}_i\}_{i=0}^n$, with $\mathcal{F}_i = \sigma(I(1), \ldots, I(i))$ - the σ-field generated by $I(1), \ldots, I(i)$ - and $\mathcal{F}_0 = \emptyset$. Clearly S_w is a sum of rank-1, independent matrices. So, if $S_w(k) = S_w - X_{I(k)}X'_{I(k)}/n$,

$$\mathbf{E}\left(\text{trace}\left([S_w^{(i)}]^{-1}\right)|\mathcal{F}_i\right) = \mathbf{E}\left(\text{trace}\left([S_w^{(i)}]^{-1}\right)|\mathcal{F}_{i-1}\right) \ .$$

The same argument as above therefore applies and the theorem is shown. □

S3.2 Proof of Theorem 3.1 and 3.2

Proof of Theorem 3.1. Recall that Wielandt's Theorem (see p.261 in [17]) gives

$$\sup_{1\leq i\leq q} \ 0 \leq \lambda_i(S_n) - \lambda_i(T_n) \leq \frac{\lambda_{max}(U_nU'_n)}{\lambda_q(T_n) - n^{-\alpha}\lambda_{max}(V_n)} \ ,$$

provided $\lambda_q(T_n) > n^{-\alpha}\lambda_{max}(V_n)$.

Recall that the Schur complement formula gives

$$n^{-\alpha}V_n \succeq U'_nT_n^{-1}U_n \succeq U'_nU_n/\lambda_{max}(T_n) \ ,$$

where the second inequality is a standard application of Lemma V.1.5 in [9]. Since $\lambda_{max}(U_nU'_n) = \lambda_{max}(U'_nU_n)$ by simply writing the singular value decomposition of U_n, we conclude that

$$\lambda_{max}(T_n)n^{-\alpha}|||V_n|||_2 \geq \lambda_{max}(U_nU'_n) \ .$$

So we conclude that provided $\lambda_q(T_n) > n^{-\alpha}\lambda_{max}(V_n)$,

$$\sup_{1\leq i\leq q} \ 0 \leq \lambda_i(S_n) - \lambda_i(T_n) \leq n^{-\alpha}\frac{\lambda_{max}(V_n)}{\lambda_q(T_n) - n^{-\alpha}\lambda_{max}(V_n)} \ . \tag{1}$$

• **Proof of Equation** (1) Note that under assumption **A2**, standard results in random matrix theory [24, 44, 45] guarantee that $|||V_n|||_2 = \mathrm{O}_P(1)$. Furthermore, standard results in classic multivariate analysis [3] show that $\lambda_q(T_n) \to \lambda_q(\Sigma_{11})$ in probability. Hence, we have

$$\frac{\lambda_{max}(V_n)}{\lambda_q(T_n) - n^{-\alpha}\lambda_{max}(V_n)} = \mathrm{O}_P(1) \ .$$

We therefore have

$$\sup_{1\leq i\leq q}\sqrt{n}(\lambda_i(S_n)-\lambda_i(T_n))=\mathrm{O}_P(n^{1/2-\alpha})\ .$$

• **Proof of Equation** (2) We note that if D_w is the diagonal matrix with the bootstrap weights on the diagonal, we have

$$S_n^*=\frac{1}{n}X'D_wX\ .$$

Therefore, we see that $|||T_n^*|||_2=\mathrm{O}_{P,w}(1)$, $\lambda_q(T_n^*)\to_{P,w}\lambda_q(\Sigma_{11})$ by the law of large numbers (provided $\mathbf{E}\,(w_i)=1$; the case of Multinomial$(n,1/n)$ weights is also easy to deal with by the technique described in the previous subsection for instance) and $|||V_n^*|||_2=\mathrm{O}_{P,w}(\mathrm{polyLog}(n))$ provided $\|w\|_\infty=\mathrm{polyLog}(n)$.

We can then conclude that

$$\sup_{1\leq i\leq q}\sqrt{n}(\lambda_i(S_n^*)-\lambda_i(T_n^*))=\mathrm{o}_{P,w}(1)\ .$$

□

Proof of Theorem 3.2. The results from Theorem 3.1 imply that

$$\sup_{1\leq i\leq q}\left|\left[\sqrt{n}(\lambda_i(S_n^*)-\lambda_i(S_n))\right]-\left[\sqrt{n}(\lambda_i(T_n^*)-\lambda_i(T_n))\right]\right|=\mathrm{o}_{P,w}(1)\ . \tag{2}$$

The arguments used in the proof of Theorem 3.1 also apply to Σ and show that

$$\sup_{1\leq i\leq q}|\lambda_i(\Sigma_n)-\lambda_i(\Sigma_{11})|\leq n^{-\alpha}\frac{\lambda_{max}(\Sigma_{22})\lambda_{max}(\Sigma_{11})}{\lambda_q(\Sigma_{11})-n^{-\alpha}\lambda_{max}(\Sigma_{22})}$$

Hence, when $\alpha>1/2+\epsilon$, we have

$$\sqrt{n}\sup_{1\leq i\leq q}|\lambda_i(\Sigma_n)-\lambda_i(\Sigma_{11})|=\mathrm{o}(1)\ .$$

Therefore,

$$\sup_{1\leq i\leq q}\sqrt{n}\,|\lambda_i(S_n)-\lambda_i(\Sigma_n)-[\lambda_i(T_n)-\lambda_i(\Sigma_{11})]|=\mathrm{o}_P(1)\ .$$

Hence, the q largest eigenvalues of S_n have the same limiting fluctuation behavior as the q largest eigenvalues of T_n (classical results [3] show that $\sqrt{n}$ is the correct order of fluctuations). The same is true for the bootstrapped version of their distributions, according to Equation (2).

Since the results of [7, 8, 17] show consistency of the bootstrap distribution of the eigenvalues of T_n, this result carries over to the bootstrap distribution of S_n. □

SUPPLEMENTARY FIGURES

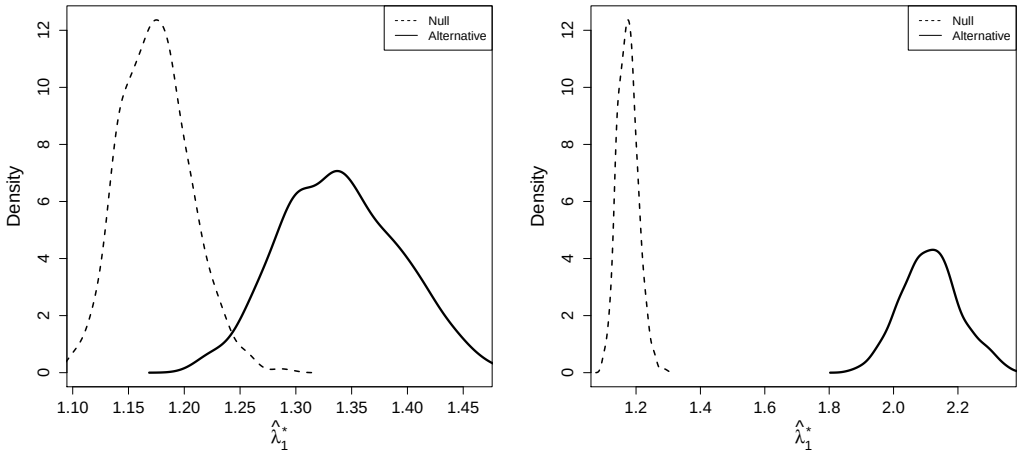

(a) $\lambda_1^{Alt} = 1 + 3\sqrt{r}$, r=0.01

(b) $\lambda_1^{Alt} = 1 + 11\sqrt{r}$, r=0.01

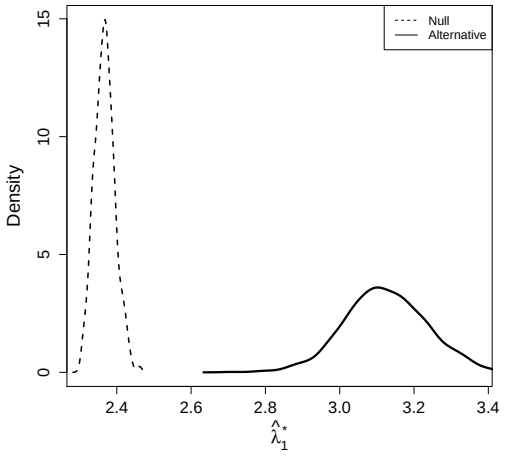

(c) $\lambda_1^{Alt} = 1 + 3\sqrt{r}$, r=0.03

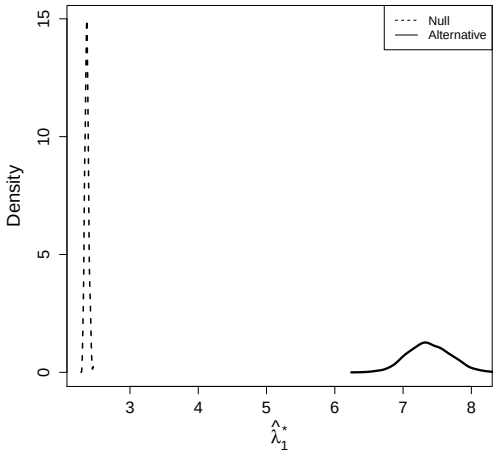

(d) $\lambda_1^{Alt} = 1 + 11\sqrt{r}$, r=0.3

Figure S1: **Top Eigenvalue: Distribution of Largest Eigenvalue, Null versus Alternative,** $X_i \sim$ **Normal**

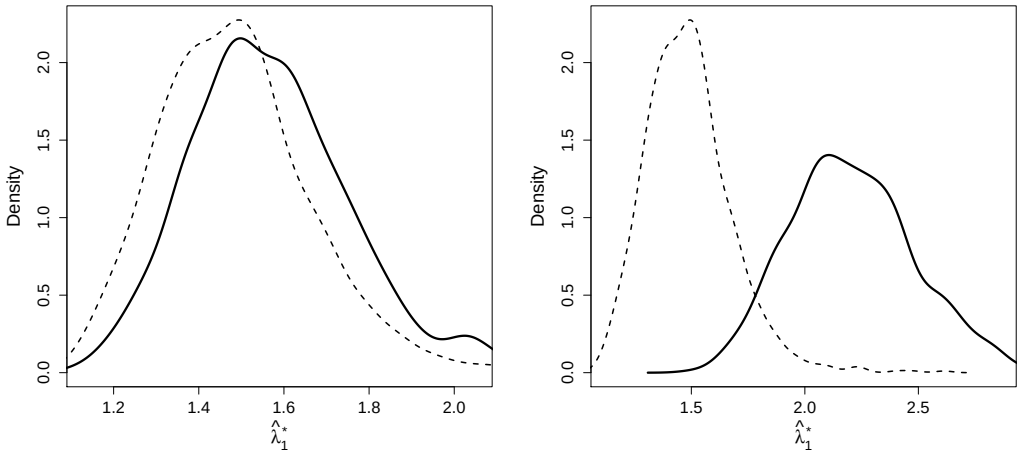

(a) $\lambda_1^{Alt} = 1 + 3\sqrt{r}$, r=0.01

(b) $\lambda_1^{Alt} = 1 + 11\sqrt{r}$, r=0.01

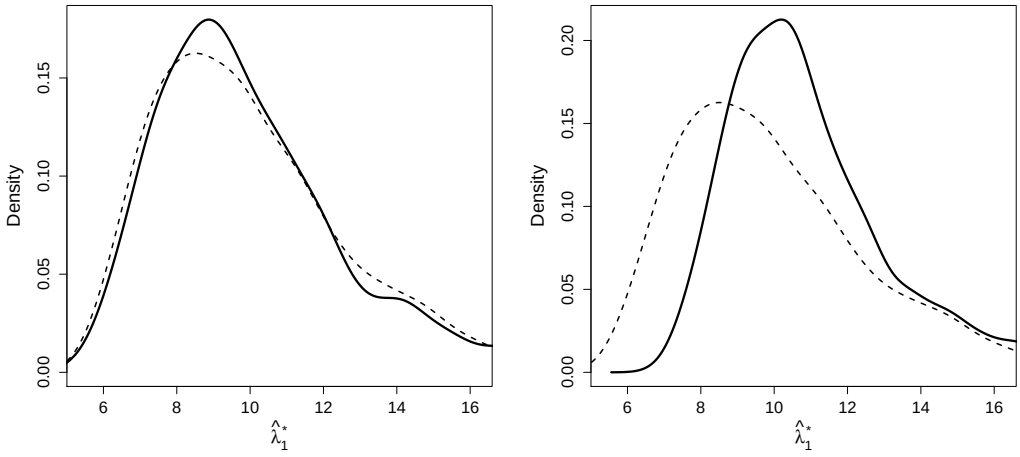

(c) $\lambda_1^{Alt} = 1 + 3\sqrt{r}$, r=0.03

(d) $\lambda_1^{Alt} = 1 + 11\sqrt{r}$, r=0.3

Figure S2: **Top Eigenvalue: Distribution of Largest Eigenvalue, Null versus Alternative,** $X_i \sim$ **Ellip. Exp**

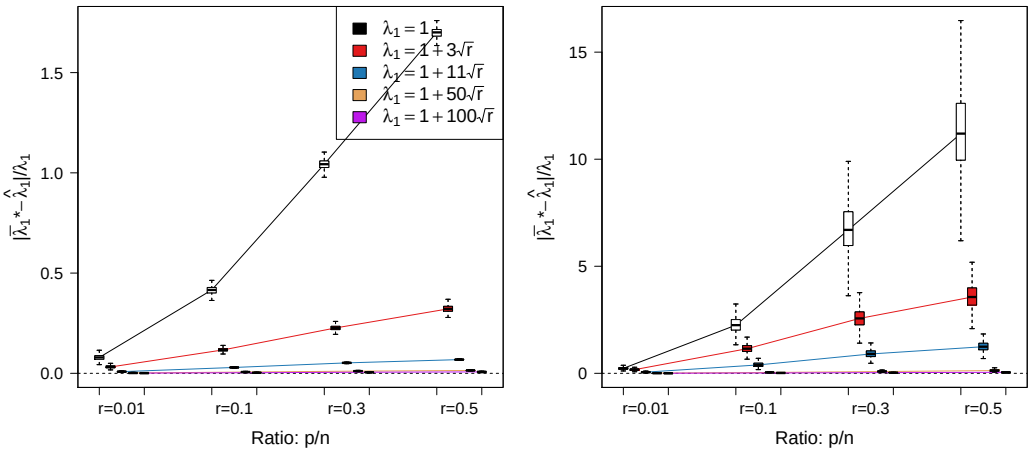

(a) $X_i \sim$ Normal

(b) $X_i \sim$ Ellip. Exp

(c) $X_i \sim$ Ellip. Normal

(d) $X_i \sim$ Ellip. Uniform

Figure S3: **Top Eigenvalue: Bias of Largest Bootstrap Eigenvalue, n=1,000:** Plotted are boxplots of the bootstrap estimate of bias covering larger values of λ_1 than shown in the main text. Unlike the main text, here we scale the bias by the true λ_1 so as to make the comparisons more comparable ($|\bar{\lambda}_1^* - \hat{\lambda}_1|/\lambda_1$). See the legend of Figure 1 in the main text for more details.

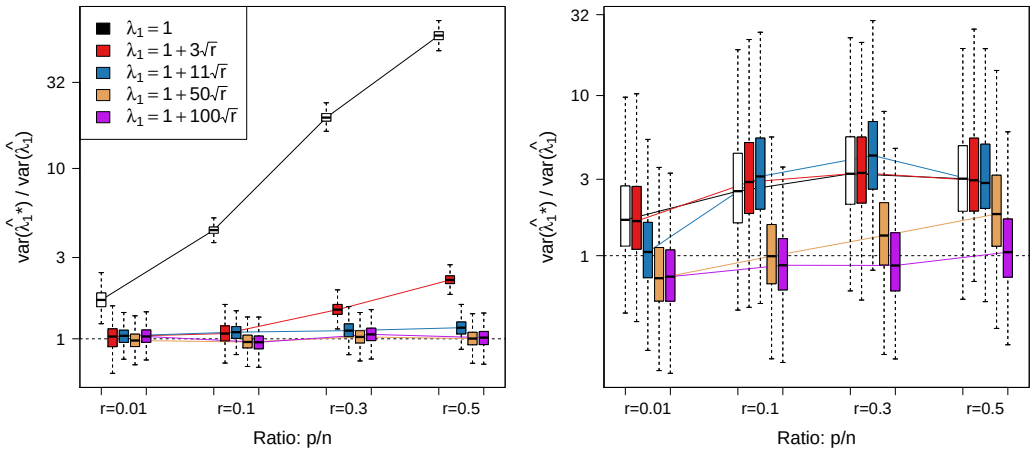

(a) $X_i \sim$ Normal

(b) $X_i \sim$ Ellip. Exp

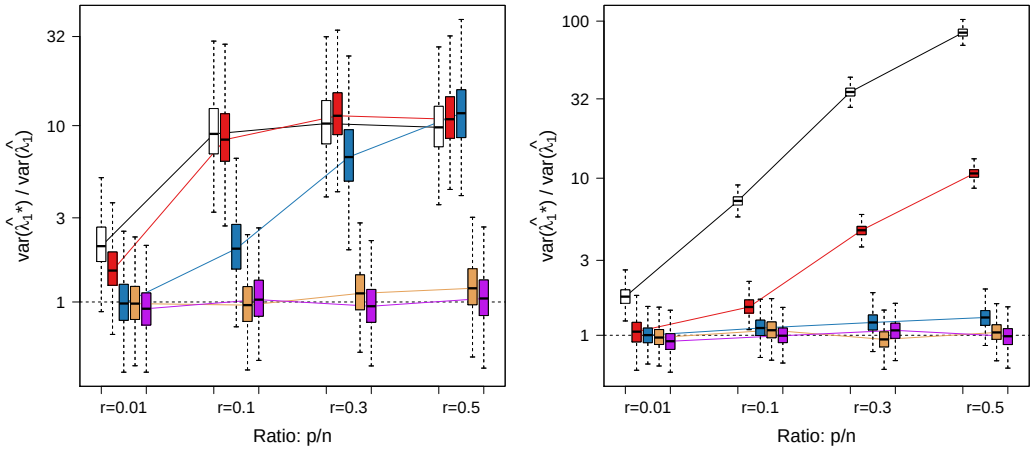

(c) $X_i \sim$ Ellip Normal

(d) $X_i \sim$ Ellip. Uniform

Figure S4: **Top Eigenvalue: Ratio of Bootstrap Estimate of Variance to True Variance for Largest Eigenvalue, n=1,000:** Plotted are boxplots of the bootstrap estimate of variance, showing larger values of λ_1 than shown in the main text. See the legend of Figure 2 in the main text for more details.

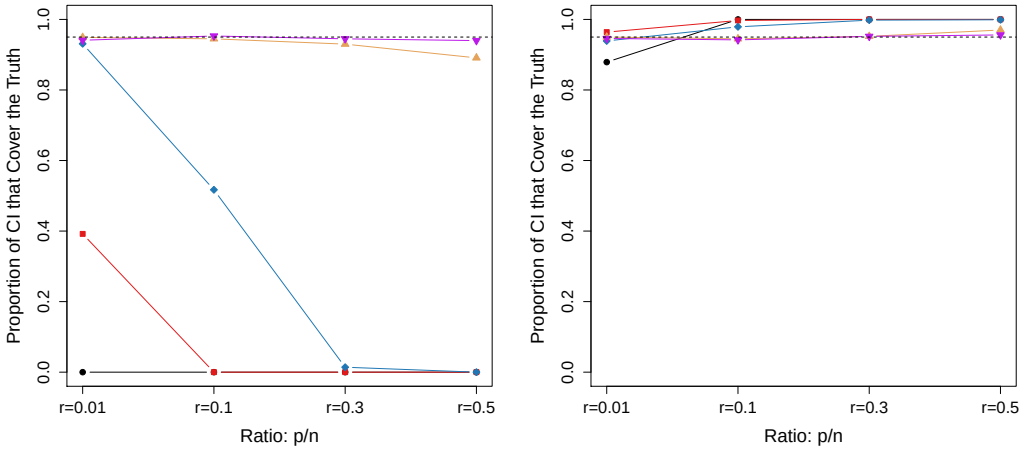

(a) $X_i \sim$ Ellip Normal
Percentile Intervals

(b) $X_i \sim$ Ellip Normal
Normal-based Intervals

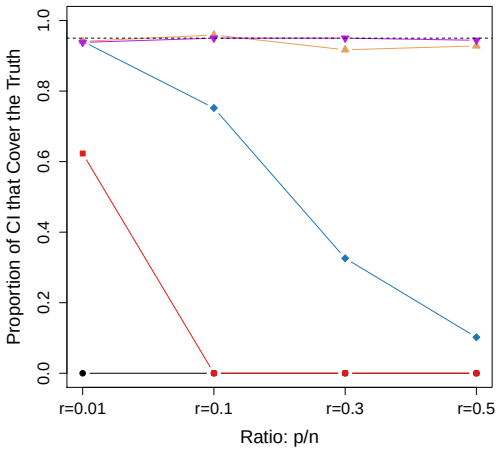

(c) $X_i \sim$ Ellip Uniform
Percentile Intervals

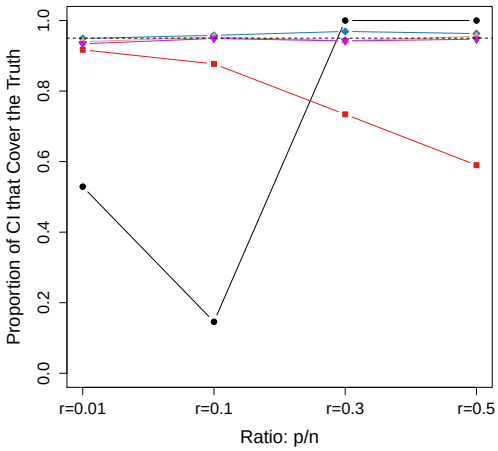

(d) $X_i \sim$ Ellip Uniform
Normal-based Intervals

Figure S5: **Top Eigenvalue: 95% CI Coverage,** $n = 1,000$ **for additional distributions:** Plotted are the corresponding CI Coverage plots for when X_i follows an elliptical distribution with Normal and Uniform weights. See Figure 4 for more details.

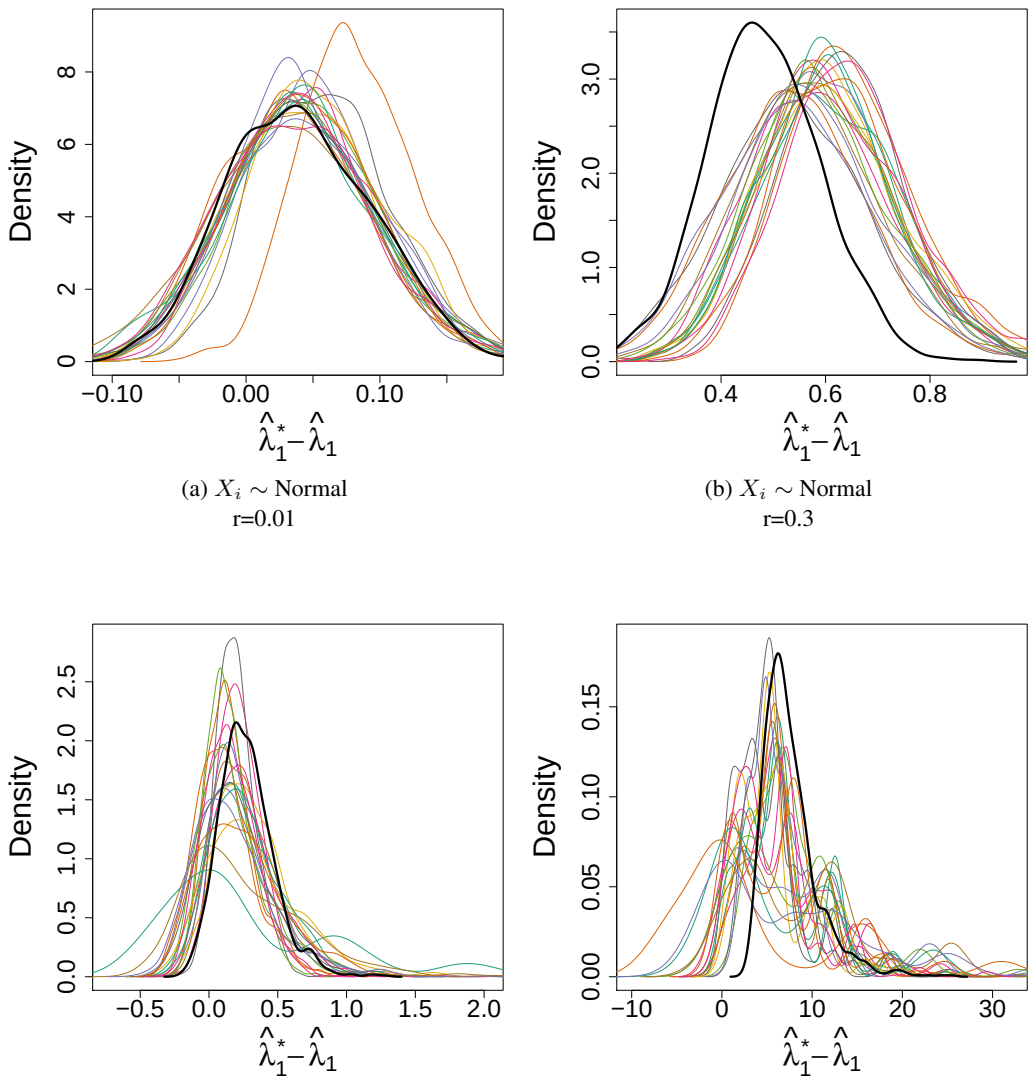

(a) $X_i \sim$ Normal r=0.01

(b) $X_i \sim$ Normal r=0.3

(c) $X_i \sim$ Ellip. Exp r=0.01

(d) $X_i \sim$ Ellip. Exp r=0.3

Figure S6: **Top Eigenvalue: Bootstrap distribution of $\hat{\lambda}_1^*$ when $\lambda_1 = 1 + 3\sqrt{r}$, n=1,000:** Plotted are the estimated density of twenty simulations of the bootstrap distribution of $\hat{\lambda}_1^{*b} - \hat{\lambda}_1$, with $b = 1, \ldots, 999$. The solid black line line represents the distribution of $\hat{\lambda}_1 - \lambda_1$ over 1,000 simulations.

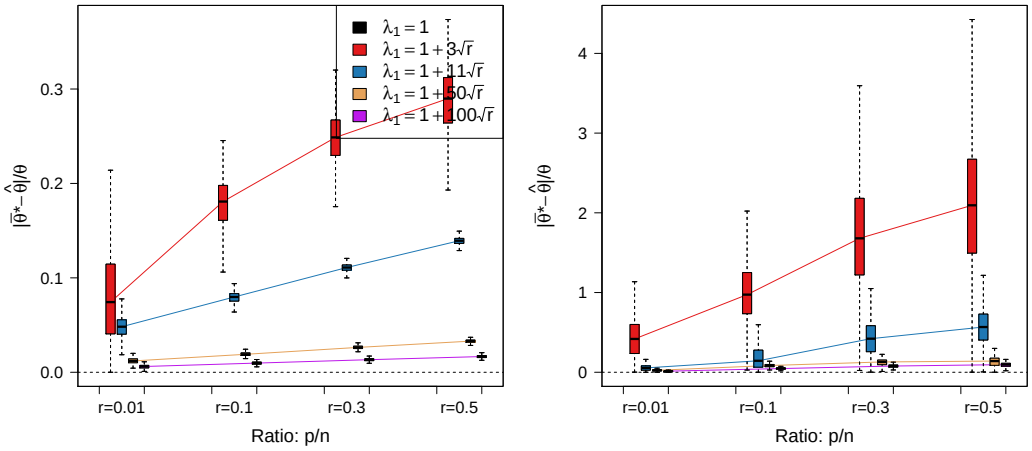

(a) $X_i \sim$ Normal

(b) $X_i \sim$ Ellip. Exp

(c) $X_i \sim$ Ellip. Norm

(d) $X_i \sim$ Ellip. Unif

Figure S7: **Gap Statistic: Ratio of Bias of Bootstrap to true Gap Statistic.** Note that the true Gap Statistic when $\lambda_1 = 1$ is zero, so that the ratio is not well defined and hence not plotted . See also Supplementary Figures S8-S11 for the median bias values and the legend of Figure 1 in the main text for more information about this plot. Note that the y-axis is different for each of the distributions, and differs from that of Supplementary Figure S3.

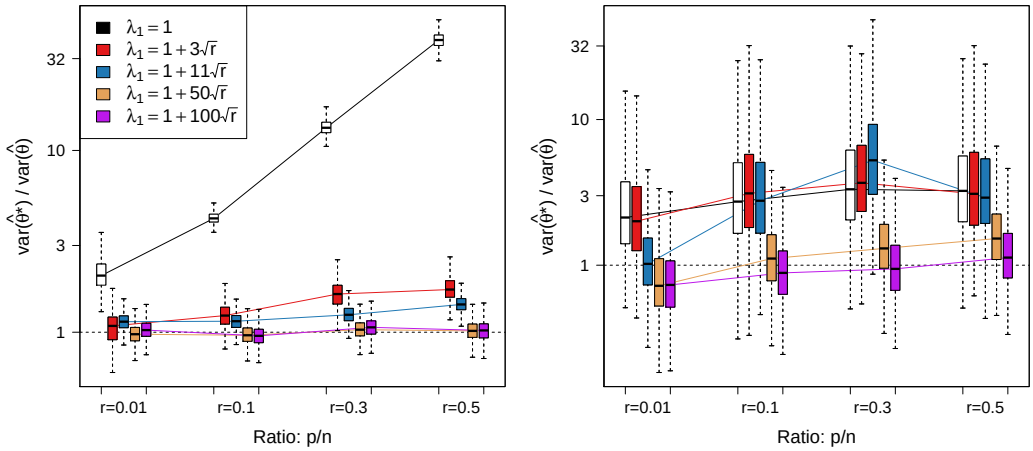

(a) $X_i \sim$ Normal

(b) $X_i \sim$ Ellip. Exp

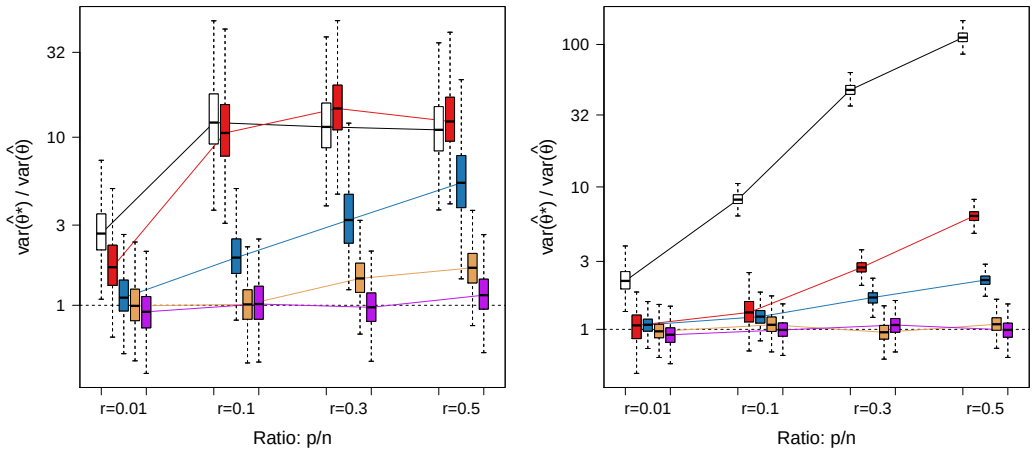

(c) $X_i \sim$ Ellip. Normal

(d) $X_i \sim$ Ellip. Uniform

Figure S8: **Gap Statistic: Ratio of Bootstrap Estimate of Variance to True Variance.**

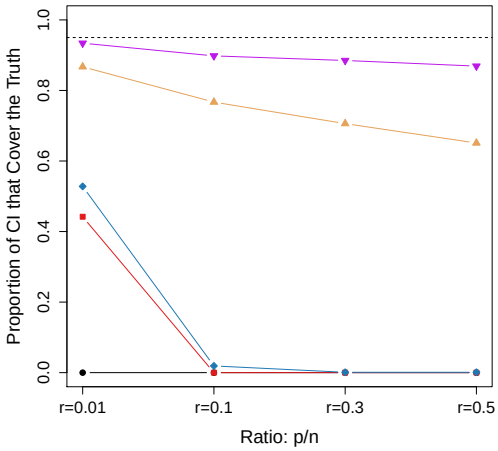

(a) $X_i \sim$ Normal
Percentile Intervals

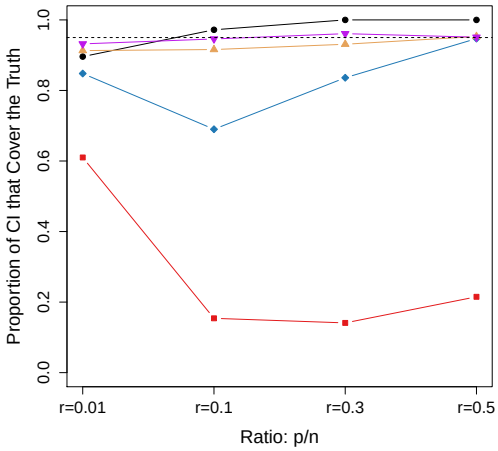

(b) $X_i \sim$ Normal
Normal-based Intervals

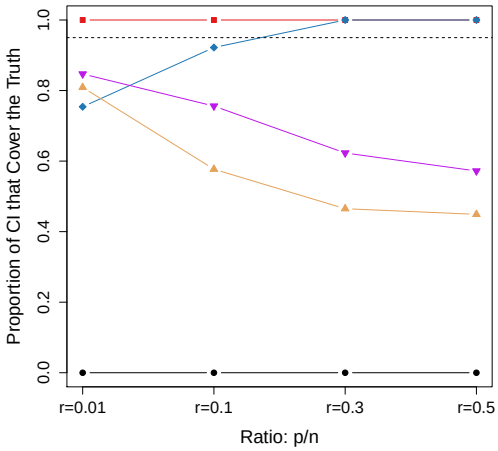

(c) $X_i \sim$ Ellip Exp
Percentile Intervals

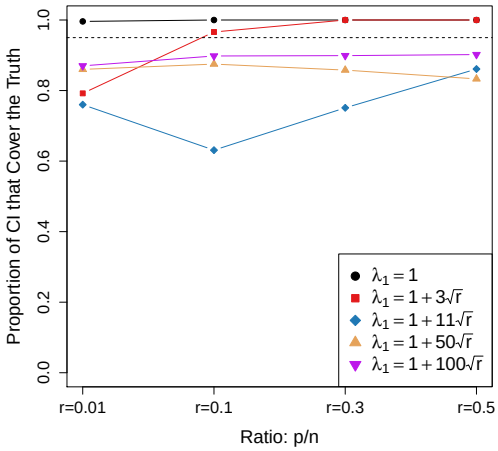

(d) $X_i \sim$ Ellip Exp
Normal-based Intervals

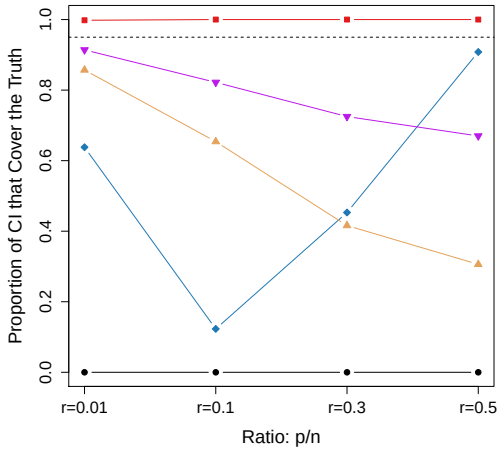

(e) $X_i \sim$ Ellip Normal
Percentile Intervals

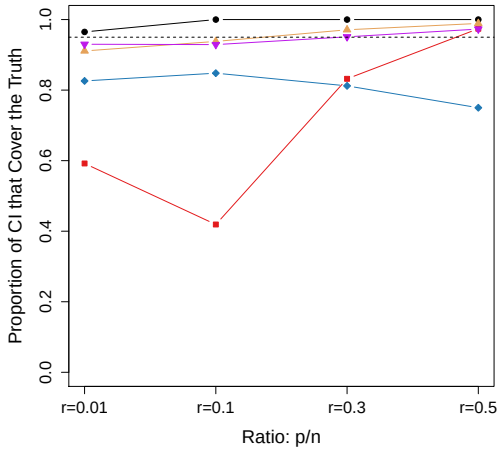

(f) $X_i \sim$ Ellip Normal
Normal-based Intervals

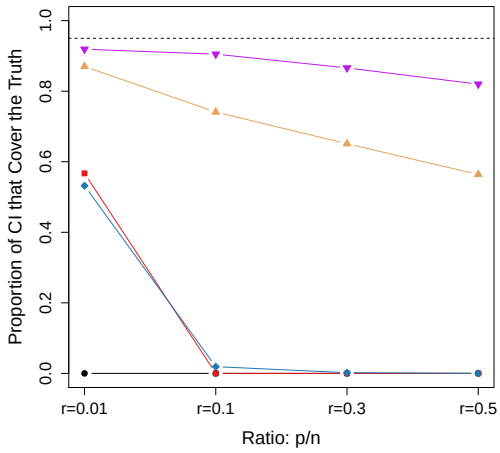

(g) $X_i \sim$ Ellip Uniform
Percentile Intervals

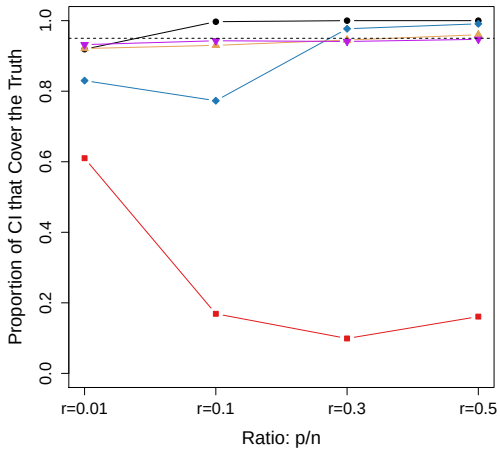

(h) $X_i \sim$ Ellip Uniform
Normal-based Intervals

Figure S8: **Gap Statistic: 95% CI Coverage,** $n = 1,000$**:**

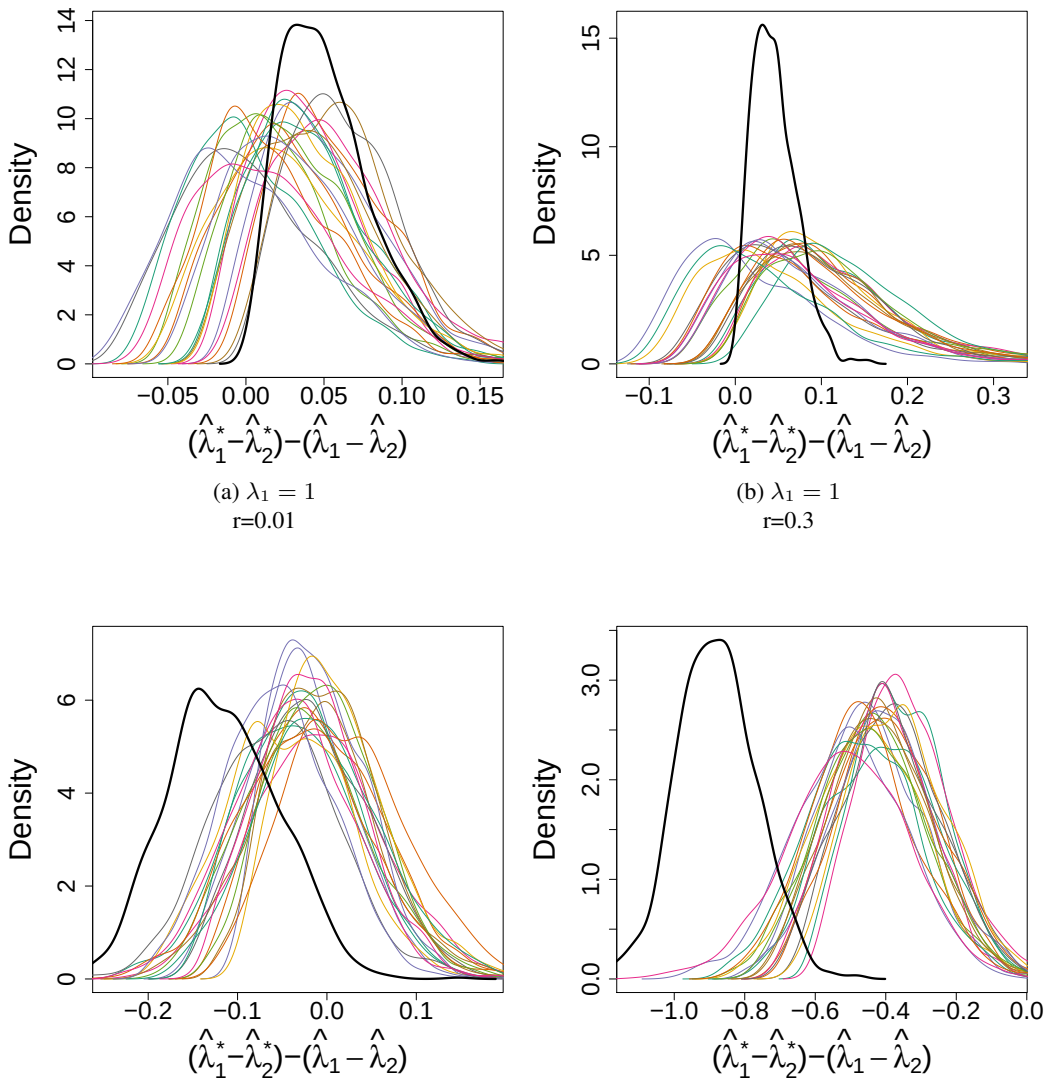

(a) $\lambda_1 = 1$
r=0.01

(b) $\lambda_1 = 1$
r=0.3

(c) $\lambda_1 = 1 + 3\sqrt{r}$
r=0.01

(d) $\lambda_1 = 1 + 3\sqrt{r}$
r=0.3

Figure S9: **Gap Statistic: Bootstrap distribution, $X_i \sim$ Normal, n=1,000:** Plotted are the estimated density of twenty simulations of the bootstrap distribution of $(\hat{\lambda}_1^{*b} - \hat{\lambda}_2^{*b}) - (\hat{\lambda}_1 - \hat{\lambda}_2)$, with $b = 1, \ldots, 999$. The solid black line line represents the distribution of $(\hat{\lambda}_1 - \hat{\lambda}_2) - (\lambda_1 - \lambda_2)$ over 1,000 simulations.

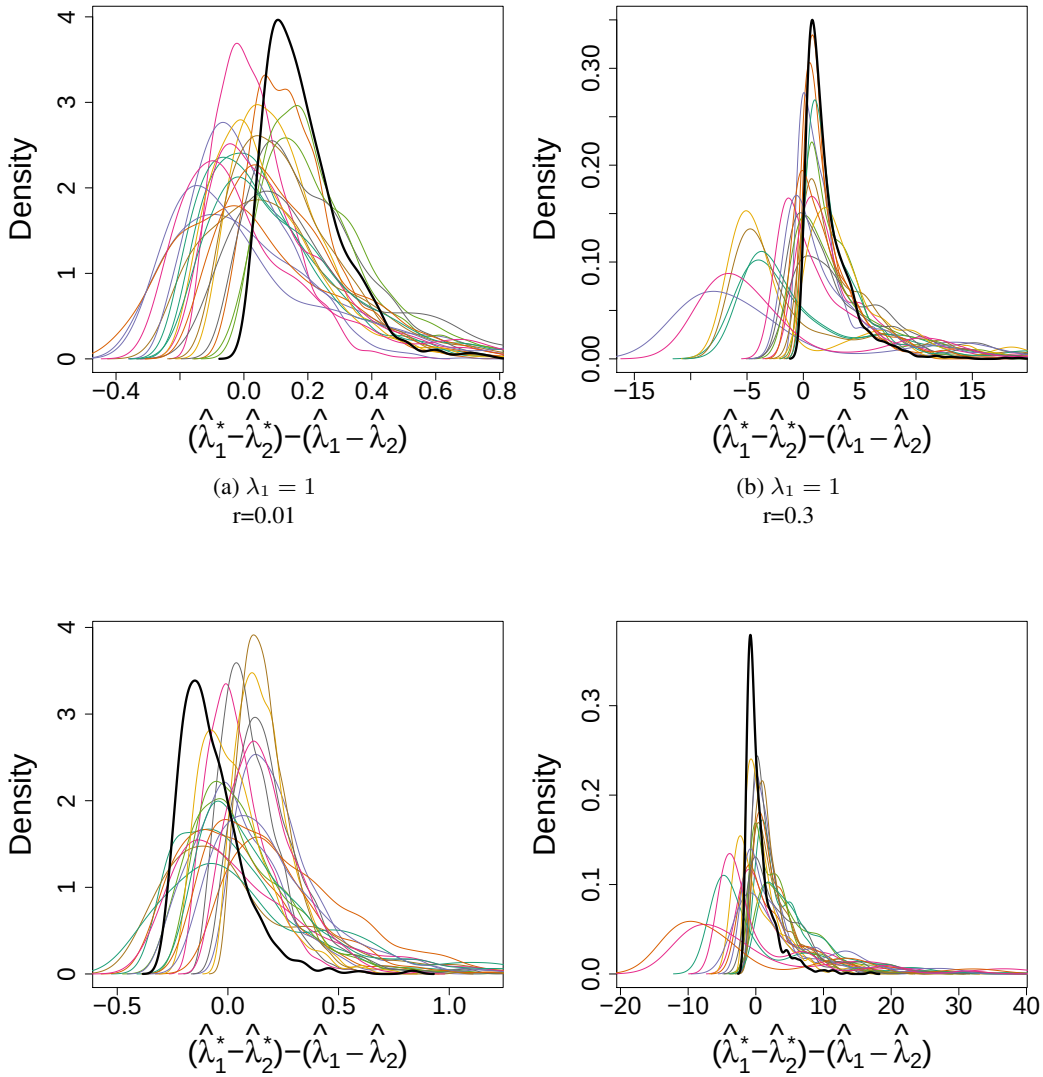

(a) $\lambda_1 = 1$
r=0.01

(b) $\lambda_1 = 1$
r=0.3

(c) $\lambda_1 = 1 + 3\sqrt{r}$
r=0.01

(d) $\lambda_1 = 1 + 3\sqrt{r}$
r=0.3

Figure S10: **Gap Statistic: Bootstrap distribution, $X_i \sim$ Ellip Exp, n=1,000:** Plotted are the estimated density of twenty simulations of the bootstrap distribution of $(\hat{\lambda}_1^{*b} - \hat{\lambda}_2^{*b}) - (\hat{\lambda}_1 - \hat{\lambda}_2)$, with $b = 1, \ldots, 999$. The solid black line line represents the distribution of $(\hat{\lambda}_1 - \hat{\lambda}_2) - (\lambda_1 - \lambda_2)$ over 1,000 simulations.

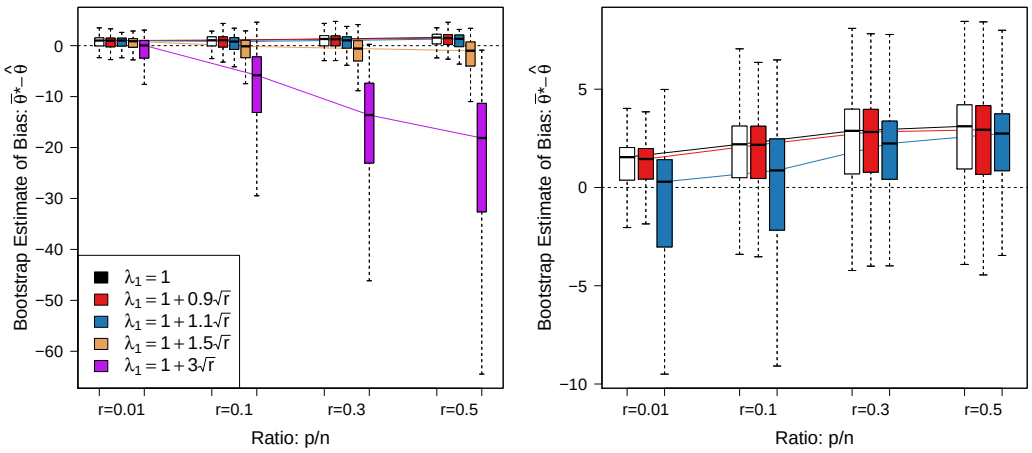

(a) $X_i \sim$ Normal

(b) $X_i \sim$ Ellip. Exp

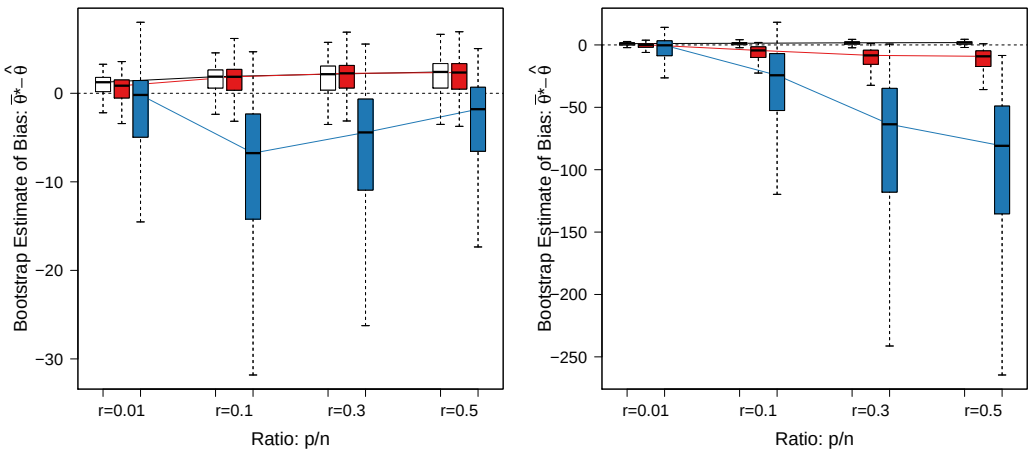

(c) $X_i \sim$ Ellip. Norm

(d) $X_i \sim$ Ellip. Unif

Figure S11: **Gap Ratio Statistic: Bias of Bootstrap.** Note that the Gap Ratio is not well defined in the population for this simulation (since $\lambda_2 = \lambda_3$) so we can not scale the Gap Ratio by the true value of the Gap Ratio as was done for other plots in the Supplementary Figures. Instead we plot the actual bias, as in Figure 1. For this reason, we only show the smaller values of λ_1 (otherwise, without scaling, the plot is dominated by the bias of large values of λ_1, even though the relative value of the bias is small). Similarly, the true bias of the estimate $\hat{\lambda}_1 - \hat{\lambda}_2$ is not a well-defined quantity and hence is not plotted.

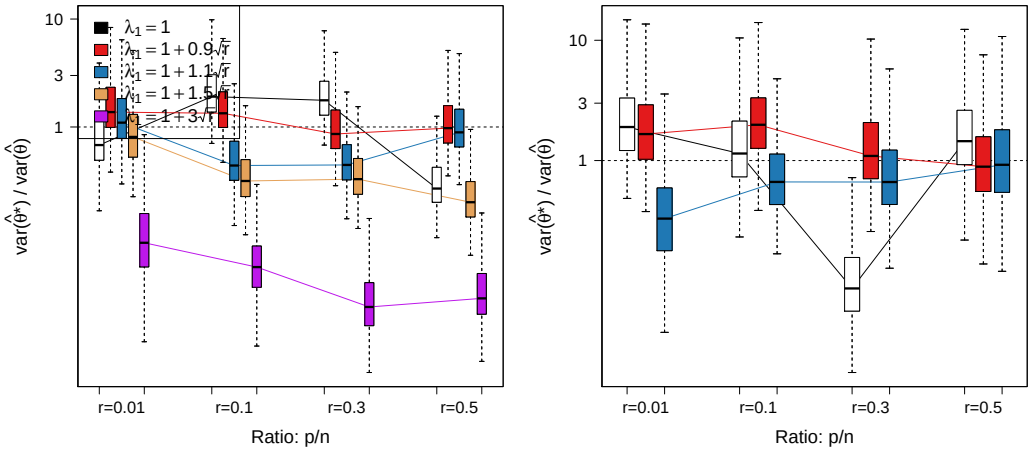

(a) $X_i \sim$ Normal

(b) $X_i \sim$ Ellip. Exp

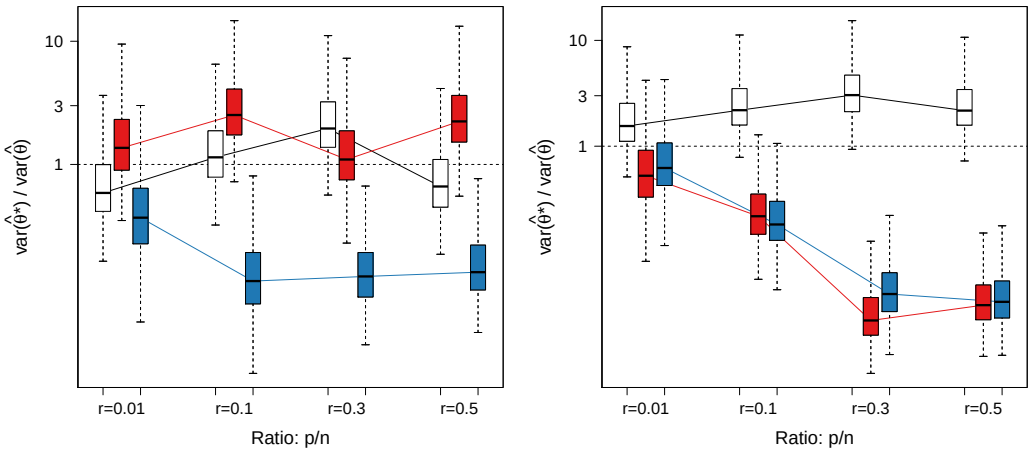

(c) $X_i \sim$ Ellip. Normal

(d) $X_i \sim$ Ellip. Uniform

Figure S12: **Gap Ratio Statistic: Ratio of Bootstrap Estimate of Variance to True Variance.**

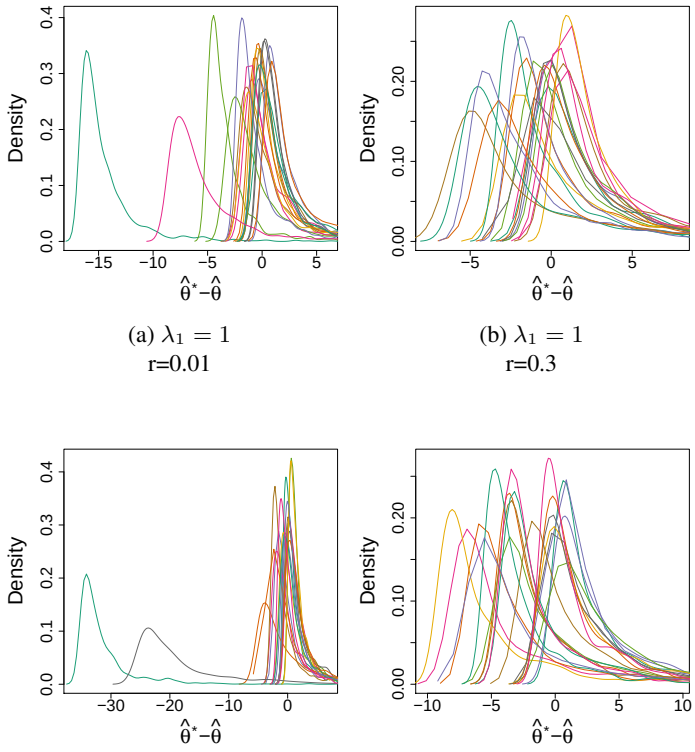

(a) $\lambda_1 = 1$
r=0.01

(b) $\lambda_1 = 1$
r=0.3

(c) $\lambda_1 = 1 + 3\sqrt{r}$
r=0.01

(d) $\lambda_1 = 1 + 3\sqrt{r}$
r=0.3

Figure S13: **Gap Ratio Statistic: Bootstrap distribution, $X_i \sim$ Ellip Exp, n=1,000:** Plotted are the estimated density of twenty simulations of the bootstrap distribution of $(\hat{\lambda}_1^{*b} - \hat{\lambda}_2^{*b}) - (\hat{\lambda}_1 - \hat{\lambda}_2)$, with $b = 1, \ldots, 999$. The solid black line line represents the distribution of $(\hat{\lambda}_1 - \hat{\lambda}_2) - (\lambda_1 - \lambda_2)$ over 1,000 simulations.

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.08	0.04	0.02	0.01	0.01
$r = 0.1$	0.42	0.23	0.13	0.11	0.10
$r = 0.3$	1.04	0.60	0.37	0.31	0.31
$r = 0.5$	1.70	1.00	0.60	0.52	0.51

(a) Bootstrap Median Estimate of Bias

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.17	0.04	0.02	0.02	0.02
$r = 0.1$	0.70	0.20	0.12	0.12	0.12
$r = 0.3$	1.37	0.48	0.36	0.28	0.38
$r = 0.5$	1.88	0.73	0.55	0.48	0.45

(b) True Bias

Table S1: **Top Eigenvalue: Median value of Bootstrap and True values of Bias, $Z \sim$ Normal** This tables give the median values of the boxplots plotted in Figure 1, as well as the true bias values (*) in the plots. See figure caption for more details.

SUPPLEMENTARY TABLES

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.23	0.22	0.12	0.06	0.05
$r = 0.1$	2.25	2.24	1.75	0.67	0.58
$r = 0.3$	6.70	6.78	6.35	2.37	1.88
$r = 0.5$	11.19	11.12	10.91	4.70	3.37

(a) Bootstrap Median Estimate of Bias

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.49	0.27	0.12	0.06	0.01
$r = 0.1$	3.34	2.39	1.10	0.54	0.75
$r = 0.3$	9.17	7.57	4.10	1.92	1.69
$r = 0.5$	14.93	12.76	8.08	3.69	3.45

(b) True Bias

Table S2: **Top Eigenvalue: Median value of Bootstrap and True values of Bias, $Z \sim$ Ellip Exp** This tables give the median values of the boxplots plotted in Figure 1, as well as the true bias values (*) in the plots. See figure caption for more details.

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.16	0.13	0.05	0.03	0.03
$r = 0.1$	1.38	1.33	0.55	0.32	0.30
$r = 0.3$	4.17	4.16	2.30	1.00	0.93
$r = 0.5$	6.95	6.96	4.80	1.72	1.57

(a) Bootstrap Median Estimate of Bias

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.32	0.13	0.05	0.04	0.07
$r = 0.1$	1.65	0.83	0.41	0.35	0.13
$r = 0.3$	4.14	2.52	1.20	0.78	0.92
$r = 0.5$	6.57	4.45	2.01	1.64	1.76

(b) True Bias

Table S3: **Top Eigenvalue: Median value of Bootstrap and True values of Bias, $Z \sim$ Ellip Norm**

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.09	0.05	0.02	0.01	0.01
$r = 0.1$	0.55	0.34	0.17	0.14	0.13
$r = 0.3$	1.56	1.07	0.49	0.40	0.39
$r = 0.5$	2.64	1.96	0.82	0.67	0.65

(a) Bootstrap Median Estimate of Bias

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.20	0.05	0.02	0.00	0.02
$r = 0.1$	0.83	0.27	0.16	0.12	0.22
$r = 0.3$	1.66	0.62	0.45	0.45	0.28
$r = 0.5$	2.33	0.98	0.72	0.63	0.63

(b) True Bias

Table S4: **Top Eigenvalue: Median value of Bootstrap and True values of Bias, $Z \sim$ Ellip Uniform**

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	1.69	1.03	1.04	0.98	1.03
$r = 0.1$	4.35	1.07	1.09	0.96	0.95
$r = 0.3$	19.88	1.48	1.12	1.03	1.06
$r = 0.5$	60.27	2.21	1.16	1.00	1.02

(a) $Z \sim$ Normal

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	1.67	1.64	1.05	0.72	0.74
$r = 0.1$	2.53	2.88	3.13	0.99	0.87
$r = 0.3$	3.25	3.29	4.22	1.34	0.87
$r = 0.5$	3.02	2.96	2.84	1.82	1.05

(b) $Z \sim$ Ellip Exp

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	2.07	1.51	0.98	0.98	0.92
$r = 0.1$	8.98	8.33	2.01	0.96	1.03
$r = 0.3$	10.27	11.37	6.63	1.12	0.95
$r = 0.5$	9.78	10.87	11.76	1.20	1.05

(c) $Z \sim$ Ellip Norm

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	1.76	1.05	1.00	0.97	0.92
$r = 0.1$	7.17	1.51	1.11	1.08	0.99
$r = 0.3$	35.38	4.66	1.21	0.94	1.07
$r = 0.5$	84.48	10.72	1.30	1.04	0.99

(d) $Z \sim$ Ellip Uniform

Table S5: **Top Eigenvalue: Median value of ratio of bootstrap estimate of variance to true variance for** $n = 1000$ This tables give the median values of the boxplots plotted in Figure 2.

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.71	0.94	0.94	0.95	0.36	0.92	0.95	0.94	0.95
$r = 0.1$	0.00	0.00	0.78	0.94	0.95	0.00	0.89	0.96	0.94	0.95
$r = 0.3$	0.00	0.00	0.39	0.94	0.96	0.04	0.84	0.96	0.95	0.96
$r = 0.5$	0.00	0.00	0.16	0.90	0.94	1.00	0.73	0.96	0.96	0.95

(a) $Z \sim$ Normal

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.30	0.94	0.92	0.92	0.99	0.99	0.95	0.91	0.90
$r = 0.1$	0.00	0.00	0.20	0.95	0.94	1.00	1.00	1.00	0.94	0.93
$r = 0.3$	0.00	0.00	0.00	0.90	0.94	1.00	1.00	1.00	0.95	0.94
$r = 0.5$	0.00	0.00	0.00	0.82	0.95	1.00	1.00	1.00	0.98	0.96

(b) $Z \sim$ Ellip Exp

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.39	0.93	0.95	0.94	0.88	0.96	0.94	0.95	0.95
$r = 0.1$	0.00	0.00	0.52	0.94	0.95	1.00	1.00	0.98	0.94	0.94
$r = 0.3$	0.00	0.00	0.01	0.93	0.94	1.00	1.00	1.00	0.95	0.95
$r = 0.5$	0.00	0.00	0.00	0.89	0.94	1.00	1.00	1.00	0.97	0.96

(c) $Z \sim$ Ellip Norm

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.62	0.94	0.94	0.94	0.53	0.92	0.95	0.94	0.93
$r = 0.1$	0.00	0.00	0.75	0.96	0.95	0.15	0.88	0.96	0.95	0.95
$r = 0.3$	0.00	0.00	0.33	0.92	0.95	1.00	0.73	0.97	0.94	0.94
$r = 0.5$	0.00	0.00	0.10	0.93	0.94	1.00	0.59	0.96	0.95	0.95

(d) $Z \sim$ Ellip Uniform

Table S6: **Top Eigenvalue: Median value of 95% CI Coverage of true λ_1 for $n = 1000$.** This tables give the percentage of CI intervals (out of 1,000 simulations) that cover the true λ_1 as plotted in Figure 4.

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.00	0.00	0.00	0.00	0.36	0.00	0.00	0.00	0.00
$r = 0.1$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$r = 0.3$	0.00	0.00	0.00	0.00	0.00	0.04	0.00	0.00	0.00	0.00
$r = 0.5$	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00

(a) $Z \sim$ Normal

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.00	0.00	0.00	0.00	0.99	0.92	0.04	0.00	0.00
$r = 0.1$	0.00	0.00	0.00	0.00	0.00	1.00	1.00	0.73	0.00	0.00
$r = 0.3$	0.00	0.00	0.00	0.00	0.00	1.00	1.00	1.00	0.01	0.00
$r = 0.5$	0.00	0.00	0.00	0.00	0.00	1.00	1.00	1.00	0.04	0.00

(b) $Z \sim$ Ellip Exp

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.00	0.00	0.00	0.00	0.88	0.47	0.00	0.00	0.00
$r = 0.1$	0.00	0.00	0.00	0.00	0.00	1.00	1.00	0.00	0.00	0.00
$r = 0.3$	0.00	0.00	0.00	0.00	0.00	1.00	1.00	0.21	0.00	0.00
$r = 0.5$	0.00	0.00	0.00	0.00	0.00	1.00	1.00	0.77	0.00	0.00

(c) $Z \sim$ Ellip Norm

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.00	0.00	0.00	0.00	0.53	0.01	0.00	0.00	0.00
$r = 0.1$	0.00	0.00	0.00	0.00	0.00	0.15	0.00	0.00	0.00	0.00
$r = 0.3$	0.00	0.00	0.00	0.00	0.00	1.00	0.01	0.00	0.00	0.00
$r = 0.5$	0.00	0.00	0.00	0.00	0.00	1.00	0.37	0.00	0.00	0.00

(d) $Z \sim$ Ellip Uniform

Table S7: **Top Eigenvalue: Median value of 95% CI Coverage of null value** 1 **for** $n = 1000$**.** This tables give the percentage of CI intervals (out of 1,000 simulations) that cover the value $\lambda_1 = 1$ for different values of the true λ_1.

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.03	-0.02	-0.05	-0.06	-0.06
$r = 0.1$	0.05	-0.17	-0.28	-0.30	-0.31
$r = 0.3$	0.10	-0.41	-0.67	-0.72	-0.73
$r = 0.5$	0.18	-0.61	-1.08	-1.17	-1.18

(a) Bootstrap Median Estimate of Bias

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.05	-0.12	-0.14	-0.14	-0.15
$r = 0.1$	0.04	-0.49	-0.58	-0.58	-0.58
$r = 0.3$	0.05	-0.88	-1.01	-1.08	-0.98
$r = 0.5$	0.05	-1.15	-1.33	-1.40	-1.43

(b) True Bias

Table S8: **Gap Statistic: Median value of Bootstrap and True values of Bias, $Z \sim$ Normal**

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.14	0.13	-0.04	-0.12	-0.14
$r = 0.1$	0.93	0.92	0.40	-1.33	-1.47
$r = 0.3$	2.74	2.76	2.54	-3.50	-4.27
$r = 0.5$	4.56	4.44	4.41	-4.82	-6.76

(a) Bootstrap Median Estimate of Bias

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.19	-0.07	-0.30	-0.37	-0.41
$r = 0.1$	0.85	-0.11	-1.88	-2.63	-2.41
$r = 0.3$	2.32	0.64	-3.51	-6.70	-7.02
$r = 0.5$	3.76	1.71	-3.74	-10.42	-11.33

(b) True Bias

Table S9: **Gap Statistic: Median value of Bootstrap and True values of Bias, $Z \sim$ Ellip Exp** This tables give the median values of the boxplots plotted in Figure 1, as well as the true bias values (*) in the plots. See figure caption for more details.

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.08	0.05	-0.07	-0.10	-0.11
$r = 0.1$	0.44	0.42	-0.66	-0.99	-1.02
$r = 0.3$	1.28	1.29	-1.00	-2.99	-3.10
$r = 0.5$	2.10	2.11	-0.34	-4.94	-5.17

(a) Bootstrap Median Estimate of Bias

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.11	-0.13	-0.24	-0.25	-0.23
$r = 0.1$	0.21	-0.70	-1.22	-1.28	-1.51
$r = 0.3$	0.56	-1.11	-2.83	-3.25	-3.13
$r = 0.5$	0.95	-1.21	-4.35	-4.80	-4.70

(b) True Bias

Table S10: **Gap Statistic: Median value of Bootstrap and True values of Bias, $Z \sim$ Ellip Norm**

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.03	-0.02	-0.06	-0.07	-0.07
$r = 0.1$	0.08	-0.17	-0.37	-0.41	-0.41
$r = 0.3$	0.24	-0.30	-1.04	-1.14	-1.15
$r = 0.5$	0.43	-0.28	-1.77	-1.94	-1.96

(a) Bootstrap Median Estimate of Bias

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	0.06	-0.12	-0.16	-0.17	-0.17
$r = 0.1$	0.05	-0.55	-0.66	-0.70	-0.60
$r = 0.3$	0.06	-1.02	-1.20	-1.20	-1.37
$r = 0.5$	0.07	-1.34	-1.60	-1.70	-1.69

(b) True Bias

Table S11: **Gap Statistic: Median value of Bootstrap and True values of Bias, $Z \sim$ Ellip Uniform**

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	2.05	1.08	1.14	0.98	1.03
$r = 0.1$	4.23	1.23	1.15	0.96	0.96
$r = 0.3$	13.33	1.62	1.25	1.03	1.06
$r = 0.5$	40.50	1.72	1.42	1.02	1.02

(a) $Z \sim$ Normal

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	2.12	2.00	1.02	0.72	0.73
$r = 0.1$	2.73	3.11	2.77	1.11	0.89
$r = 0.3$	3.32	3.68	5.25	1.30	0.94
$r = 0.5$	3.24	3.09	2.91	1.52	1.13

(b) $Z \sim$ Ellip Exp

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	2.67	1.68	1.11	0.99	0.91
$r = 0.1$	12.22	10.60	1.92	1.01	1.02
$r = 0.3$	11.52	14.83	3.22	1.44	0.97
$r = 0.5$	11.07	12.41	5.36	1.67	1.15

(c) $Z \sim$ Ellip Norm

	$\lambda_1 = 1$	$\lambda_1 = 1 + 3\sqrt{r}$	$\lambda_1 = 1 + 11\sqrt{r}$	$\lambda_1 = 1 + 50\sqrt{r}$	$\lambda_1 = 1 + 100\sqrt{r}$
$r = 0.01$	2.19	1.07	1.08	0.97	0.92
$r = 0.1$	8.14	1.31	1.23	1.08	0.99
$r = 0.3$	47.97	2.72	1.67	0.95	1.07
$r = 0.5$	111.75	6.25	2.22	1.09	0.99

(d) $Z \sim$ Ellip Uniform

Table S12: **Gap Statistic: Median value of ratio of bootstrap estimate of variance to true variance for** $n = 1000$ This tables give the median values of the boxplots plotted in SupplementaryFigure S8.

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.44	0.53	0.87	0.93	0.90	0.61	0.85	0.91	0.93
$r = 0.1$	0.00	0.00	0.02	0.77	0.90	0.97	0.15	0.69	0.92	0.95
$r = 0.3$	0.00	0.00	0.00	0.71	0.89	1.00	0.14	0.84	0.93	0.96
$r = 0.5$	0.00	0.00	0.00	0.65	0.87	1.00	0.21	0.95	0.95	0.95

(a) $Z \sim$ Normal

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	1.00	0.75	0.81	0.85	1.00	0.79	0.76	0.86	0.87
$r = 0.1$	0.00	1.00	0.92	0.58	0.76	1.00	0.97	0.63	0.88	0.90
$r = 0.3$	0.00	1.00	1.00	0.47	0.62	1.00	1.00	0.75	0.86	0.90
$r = 0.5$	0.00	1.00	1.00	0.45	0.57	1.00	1.00	0.86	0.83	0.90

(b) $Z \sim$ Ellip Exp

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	1.00	0.64	0.86	0.91	0.96	0.59	0.83	0.91	0.93
$r = 0.1$	0.00	1.00	0.12	0.65	0.82	1.00	0.42	0.85	0.94	0.93
$r = 0.3$	0.00	1.00	0.45	0.42	0.72	1.00	0.83	0.81	0.97	0.95
$r = 0.5$	0.00	1.00	0.91	0.31	0.67	1.00	0.97	0.75	0.99	0.97

(c) $Z \sim$ Ellip Norm

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.57	0.53	0.87	0.92	0.92	0.61	0.83	0.92	0.93
$r = 0.1$	0.00	0.00	0.02	0.74	0.91	1.00	0.17	0.77	0.93	0.94
$r = 0.3$	0.00	0.00	0.00	0.65	0.87	1.00	0.10	0.98	0.94	0.94
$r = 0.5$	0.00	0.00	0.00	0.56	0.82	1.00	0.16	0.99	0.96	0.95

(d) $Z \sim$ Ellip Uniform

Table S13: **Gap Statistic: Median value of 95% CI Coverage of true Gap for $n = 1000$.** This tables give the percentage of CI intervals (out of 1,000 simulations) that cover the true $\lambda_1 - \lambda_2$ as plotted in Supplementary Figure S8.

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.00	0.85	0.00	0.00	0.00	0.00	0.95	0.00	0.00
$r = 0.1$	0.00	0.00	0.00	0.00	0.00	0.00	0.08	0.00	0.00	0.00
$r = 0.3$	0.00	0.00	0.00	0.00	0.00	0.00	0.78	0.00	0.00	0.00
$r = 0.5$	0.00	0.01	0.00	0.00	0.00	0.00	0.12	0.00	0.00	0.00

(a) $Z \sim$ Normal

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.28	0.33	0.87	0.00	0.00	0.07	0.08	0.84	0.00	0.00
$r = 0.1$	1.00	1.00	1.00	0.00	0.00	0.95	0.95	0.96	0.01	0.00
$r = 0.3$	1.00	1.00	0.99	0.00	0.00	1.00	1.00	1.00	0.04	0.00
$r = 0.5$	0.97	0.97	0.92	0.00	0.00	1.00	1.00	1.00	0.10	0.01

(b) $Z \sim$ Ellip Exp

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.00	0.79	0.00	0.00	0.00	0.00	0.92	0.00	0.00
$r = 0.1$	1.00	1.00	0.96	0.00	0.00	0.27	0.35	0.07	0.00	0.00
$r = 0.3$	1.00	1.00	0.99	0.00	0.00	1.00	1.00	0.39	0.00	0.00
$r = 0.5$	1.00	1.00	1.00	0.00	0.00	1.00	1.00	0.78	0.00	0.00

(c) $Z \sim$ Ellip Norm

	Percentile					Normal				
	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$	$\lambda_1 = 1$	$\lambda_1 = 3\sqrt{r}$	$\lambda_1 = 11\sqrt{r}$	$\lambda_1 = 50\sqrt{r}$	$\lambda_1 = 100\sqrt{r}$
$r = 0.01$	0.00	0.00	0.83	0.00	0.00	0.00	0.00	0.94	0.00	0.00
$r = 0.1$	0.00	0.00	0.00	0.00	0.00	0.00	0.10	0.00	0.00	0.00
$r = 0.3$	0.07	0.03	0.00	0.00	0.00	0.00	0.90	0.00	0.00	0.00
$r = 0.5$	1.00	1.00	0.00	0.00	0.00	0.00	0.98	0.00	0.00	0.00

(d) $Z \sim$ Ellip Uniform

Table S14: **Gap Statistic: Median value of 95% CI Coverage of null value** 0 **for** $n = 1000$**.** This tables give the percentage of CI intervals (out of 1,000 simulations) that cover the value $\lambda_1 - \lambda_2 = 0$ for different values of the true λ_1.

References

[1] Demissie Alemayehu. Bootstrapping the latent roots of certain random matrices. *Comm. Statist. Simulation Comput.*, 17(3):857–869, 1988.

[2] T. W. Anderson. Asymptotic theory for principal component analysis. *Ann. Math. Statist.*, 34:122–148, 1963.

[3] T. W. Anderson. *An introduction to multivariate statistical analysis*. Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, third edition, 2003.

[4] Hamid Babamoradi, Frans van den Berg, and Åsmund Rinnan. Bootstrap based confidence limits in principal component analysis – a case study. *Chemometrics and Intelligent Laboratory Systems*, 120:97–105, January 2013.

[5] Z. D. Bai. Methodologies in spectral analysis of large-dimensional random matrices, a review. *Statist. Sinica*, 9(3):611–677, 1999. With comments by G. J. Rodgers and Jack W. Silverstein; and a rejoinder by the author.

[6] J. Baik, G. Ben Arous, and S. Péché. Phase transition of the largest eigenvalue for non-null complex sample covariance matrices. *Ann. Probab.*, 33(5):1643–1697, 2005.

[7] Rudolf Beran and Muni S. Srivastava. Bootstrap tests and confidence regions for functions of a covariance matrix. *Ann. Statist.*, 13(1):95–115, 1985.

[8] Rudolf Beran and Muni S. Srivastava. Correction: "Bootstrap tests and confidence regions for functions of a covariance matrix" [Ann. Statist. **13** (1985), no. 1, 95–115; MR0773155 (86g:62054)]. *Ann. Statist.*, 15(1):470–471, 1987.

[9] Rajendra Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.

[10] Peter J. Bickel and David A. Freedman. Some asymptotic theory for the bootstrap. *Ann. Statist.*, 9(6):1196–1217, 1981.

[11] A. Boutet de Monvel, A. Khorunzhy, and V. Vasilchuk. Limiting eigenvalue distribution of random matrices with correlated entries. *Markov Process. Related Fields*, 2(4):607–636, 1996.

[12] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.

[13] A. C. Davison and D. V. Hinkley. *Bootstrap methods and their application*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 1997.

[14] Persi Diaconis and David Freedman. Asymptotics of graphical projection pursuit. *Ann. Statist.*, 12(3):793–815, 1984.

[15] Momar Dieng. Distribution functions for edge eigenvalues in orthogonal and symplectic ensembles: Painlevé representations. *Int. Math. Res. Not.*, 37(37):2263–2287, 2005.

[16] Lutz Dümbgen. On nondifferentiable functions and the bootstrap. *Probab. Theory Related Fields*, 95(1):125–140, 1993.

[17] Morris L. Eaton and David E. Tyler. On Wielandt's inequality and its application to the asymptotic distribution of the eigenvalues of a random symmetric matrix. *Ann. Statist.*, 19(1):260–271, 1991.

[18] B. Efron. Bootstrap methods: another look at the jackknife. *Ann. Statist.*, 7(1):1–26, 1979.

[19] Noureddine El Karoui. On the largest eigenvalue of Wishart matrices with identity covariance when n, p and $p/n \to \infty$. *arXiv:math.ST/0309355*, September 2003.

[20] Noureddine El Karoui. Tracy-Widom limit for the largest eigenvalue of a large class of complex sample covariance matrices. *The Annals of Probability*, 35(2):663–714, March 2007.

[21] Noureddine El Karoui. Concentration of measure and spectra of random matrices: Applications to correlation matrices, elliptical distributions and beyond. *The Annals of Applied Probability*, 19(6):2362–2405, December 2009.

[22] Noureddine El Karoui and Elizabeth Purdom. Can we trust the bootstrap in high-dimension? Technical Report 824, UC Berkeley, Department of Statistics, February 2015. Submitted to AoS.

[23] Aaron Fisher, Brian Caffo, Brian Schwartz, and Vadim Zipunnikov. Fast, exact bootstrap principal component analysis for $p > 1$ million. *Journal of the American Statistical Association*, 2015.

[24] S. Geman. A limit theorem for the norm of random matrices. *Ann. Probab.*, 8(2):252–261, 1980.

[25] Jeffrey S. Geronimo and Theodore P. Hill. Necessary and sufficient condition that the limit of Stieltjes transforms is a Stieltjes transform. *J. Approx. Theory*, 121(1):54–60, 2003.

[26] F. Götze and A. N. Tikhomirov. Limit theorems for spectra of random matrices with martingale structure. In *Stein's method and applications*, volume 5 of *Lect. Notes Ser. Inst. Math. Sci. Natl. Univ. Singap.*, pages 181–193. Singapore Univ. Press, Singapore, 2005.

[27] Peter Hall. *The bootstrap and Edgeworth expansion*. Springer Series in Statistics. Springer-Verlag, New York, 1992.

[28] Peter Hall, Young K. Lee, Byeong U. Park, and Debashis Paul. Tie-respecting bootstrap methods for estimating distributions of sets and functions of eigenvalues. *Bernoulli*, 15(2):380–401, 2009.

[29] Peter Hall, J. S. Marron, and Amnon Neeman. Geometric representation of high dimension, low sample size data. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(3):427–444, 2005.

[30] Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, 1990. Corrected reprint of the 1985 original.

[31] Iain M. Johnstone. High dimensional statistical inference and random matrices. In *International Congress of Mathematicians. Vol. I*, pages 307–333. Eur. Math. Soc., Zürich, 2007.

[32] I.M. Johnstone. On the distribution of the largest eigenvalue in principal component analysis. *Ann. Statist.*, 29(2):295–327, 2001.

[33] Tosio Kato. *Perturbation theory for linear operators*. Classics in Mathematics. Springer-Verlag, Berlin, 1995. Reprint of the 1980 edition.

[34] M. Ledoux. *The concentration of measure phenomenon*, volume 89 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2001.

[35] Ji oon Lee and Kevin Schnelli. Tracy-widom distribution for the largest eigenvalue of real sample covariance matrices with general population. *arxiv:1409.4979*, 2014.

[36] V. A. Marčenko and L. A. Pastur. Distribution of eigenvalues in certain sets of random matrices. *Mat. Sb. (N.S.)*, 72 (114):507–536, 1967.

[37] Richard O. Michaud. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. Oxford University Press, USA, June 1998.

[38] A. Onatski. Talk at mit conference, July 2006.

[39] Alexei Onatski. The Tracy-Widom limit for the largest eigenvalues of singular complex Wishart matrices. *Ann. Appl. Probab.*, 18(2):470–490, 2008.

[40] A. Pajor and L. Pastur. On the limiting empirical measure of eigenvalues of the sum of rank one matrices with log-concave distribution. *Studia Math.*, 195(1):11–29, 2009.

[41] Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617–1642, October 2007.

[42] Debashis Paul and Jack W. Silverstein. No eigenvalues outside the support of the limiting empirical spectral distribution of a separable covariance matrix. *J. Multivariate Anal.*, 100(1):37–57, 2009.

[43] Dimitris N. Politis, Joseph P. Romano, and Michael Wolf. *Subsampling*. Springer Series in Statistics. Springer-Verlag, New York, 1999.

[44] Jack W. Silverstein. The smallest eigenvalue of a large-dimensional Wishart matrix. *Ann. Probab.*, 13(4):1364–1368, 1985.

[45] Jack W. Silverstein. On the weak limit of the largest eigenvalue of a large-dimensional sample covariance matrix. *J. Multivariate Anal.*, 30(2):307–311, 1989.

[46] Jack W. Silverstein. Strong convergence of the empirical distribution of eigenvalues of large-dimensional random matrices. *J. Multivariate Anal.*, 55(2):331–339, 1995.

[47] Alexander Soshnikov. Universality at the edge of the spectrum in Wigner random matrices. *Comm. Math. Phys.*, 207(3):697–733, 1999.

[48] Stephen K. Mitchell Sungjin. Hong and Richard A. Harshman. Bootstrap scree tests: A monte carlo simulation and applications to published data. *British Journal of Mathematical and Statistical Psychology*, 59(1):35–57, May 2006.

[49] M. E. Timmerman, H. A. Kiers, and A. K. Smilde. Estimating confidence intervals for principal component loadings: A comparison between the bootstrap and asymptotic results. *British Journal of Mathematical and Statistical Psychology*, 60(2):295–314, 2007.

[50] Kenneth W. Wachter. The strong limits of random matrix spectra for sample matrices of independent elements. *Annals of Probability*, 6(1):1–18, 1978.